SG24-4571-00

TME 10 Information/Management Version 1.1 Performance and Capacity Considerations

September 1997

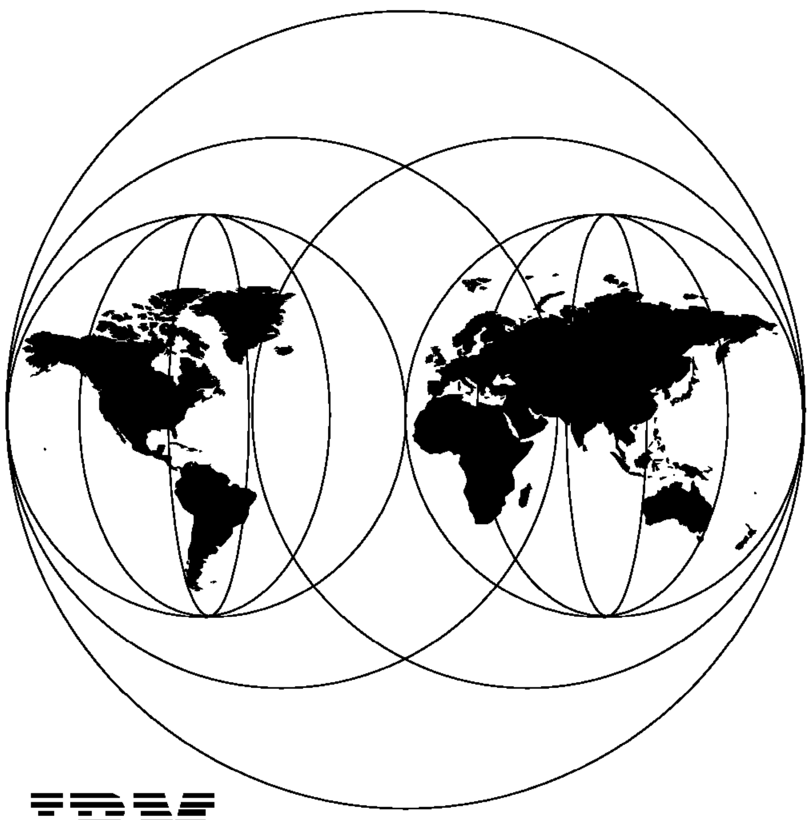

International Technical Support Organization
Raleigh Center

International Technical Support Organization

SG24-4571-00

TME 10 Information/Management Version 1.1 Performance and Capacity Considerations

September 1997

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix D, "Special Notices" on page 93.

First Edition (September 1997)

This edition applies to Version 1, Release 1 of the TME 10 Information/Management, Program Number 5648-142, for use with the MVS/ESA Operating System, and to Version 6, Release 3 of Information/Management, Program Number 5695-171, for use with the MVS/ESA Operating System.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 678
P.O. Box 12195
Research Triangle Park, NC 27709-2195

Contents

Preface

This redbook examines, explains and provides guidance when facing Information/Management performance issues. Techniques to gather and analyze performance data, and to identify and avoid performance bottlenecks are described in detail.

It will be extremely helpful to people who need to understand the tasks required to plan and tune Version 1.1 in network computing and client/server environments.

Several practical examples from a customer scenario are presented to demonstrate some of the points and recommendations made in this redbook.

It is assumed that the reader is not only familiar with the Information/Management product, but also has a good knowledge of MVS/ESA and VSAM.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the Systems Management and Networking ITSO Center, Raleigh.

Sergio Juri is an Advisory Technical Support Specialist at the Systems Management and Networking ITSO Center, Raleigh. He writes extensively on systems management with special emphasis on problem management on the enterprise intranet. He has been guest speaker at the TMC/96, AOTC/96 and German Guide conferences, and many IBM internal conferences and workshops. Before joining the ITSO three years ago, he worked with large and distributed systems in his home country, Argentina. He has been a member of Argentina′s Information Science Professional Council (CPCI, Consejo Profesional de Ciencia Informatica) since 1988. His areas of expertise include large systems, distributed systems, network computing and systems management. He has also designed and implemented different migration techniques for IBM and IBM customers worldwide, including many large customers in the US and Germany. He has also worked extensively with network computing and Internet enablement.

Francisco Gayo is a project leader and engagement manager with the Customer Solutions and Services organization in Spain. He has more than 20 years of experience in IT, including three years of experience in ISM consulting and five more years in leading, designing and implementing ISM projects. He holds a degree in Telecommunications Engineering from the ETSIT located in Madrid, Spain. He was a staff member at the ITSO Center in Raleigh from 1988 until 1992. His relevant areas of expertise include ISM assessment, design and sizing, enterprise-wide event handling systems, and IT infrastructure design. He was a team member of IBM′s group of experts who initially developed the End-to-End Systems Design method. He has written on capacity and performance planning and measurement, pro-active monitoring systems, and IT generic design guidelines and methodologies.

Helmut Graef is an Information/Management Specialist in Germany. He has over 10 years of experience with the Information/Management product. His areas of expertise include Systems Management, Problem and Change Management, and performance tuning. He has written extensively on the Information/Management product, specially in the area of performance tuning.

Don Miller is an Advisory Software Engineer in Raleigh, NC. Don joined IBM in May 1977 as a Customer Engineer (CE). He was a resident CE for the University of Auburn in Auburn, Alabama where he supported multiple mainframe computers and a network of approximately 500 terminals. After 7 years in hardware, he moved to Tampa in 1984 and started a career in software as a Level 1 Program Support Representative (PSR) at the Tampa Software Support Center. His area of specialty was MVS including JES2, JES3, and SMP/E. In 1988, Don accepted a job as a Level 2 Operational Support Specialist for Information/Management and relocated to Raleigh, North Carolina. Today, Don is involved in every aspect of the Information/Management product as an advisory software engineer, but he′s still faithful to his first love, staying in close contact with customers.

Bobby Dixon is a Senior IT Specialist for IBM South Africa. He has 13 years of experience in the IT field. His areas of expertise include Information/Management, ASIM and MVS.

Thanks to the following people for their invaluable contributions to this project:

Gail Wojton
International Technical Support Organization, Raleigh Center

Rollin Hippler
Tivoli Systems, RTP, North Carolina

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in “ITSO Redbook Evaluation” on page 107 to the fax number shown on the form.
- Use the electronic evaluation form found on the Redbooks Web sites:

For Internet users	`http://www.redbooks.ibm.com`
For IBM Intranet users	`http://w3.itso.ibm.com`

- Send us a note at the following address:

  `redbook@vnet.ibm.com`

Chapter 1. An Overview of Information/Management′s Performance Management

Performance management is essentially the art of managing a finite set of resources maximizing the overall throughput while minimizing user and/or service response times.

This involves decisions concerning which groups of users, or which transaction types, have the greatest priority. Hence, sometimes a subset of the user population, probably the batch users, has to be given less priority for the health of the system as a whole. Other users, such as those belonging to help desks, will receive better service.

1.1 READ THIS FIRST! Product Name and Release Number Change

The Information/Management product originated in the IBM company, and has been very successful for many years. After the IBM/Tivoli merger, Information/Management changed to reflect Tivoli standards, and was renamed TME 10 Information/Management.

The information in this book applies mostly to the TME 10 Information/Management Version 1.1 and Information/Management Version 6.3 products. References to Information/Management apply to both TME 10 Information/Management and Information/Management. When a distinction is necessary, the release number will be explicitly indicated.

Keep in mind that TME 10 Information/Management Version 1.1 is the release that follows Information/Management Version 6.3. The latest version of the product is then Version 1.1.

1.2 A Rule of Thumb Approach for Identifying Performance Issues

Information/Management specialists may approach performance problems differently depending upon their experience and problem solving methods. Sometimes they will be able to easily identify a problem and go straight to the appropriate solution. In other situations, a previous task of data gathering is mandatory for diagnosis purposes.

Prior to starting any data collection, you may want to use a bullet-proof and quite effective technique which would help you isolate the problem area. First of all, you better ask your customer these basic questions:

- How many users does Information/Management have? How many concurrently?
- How many records does Information/Management have in its SDDS and SDIDS databases?

The responses to the above questions may shed some light to help diagnose the problem. Everything might depend on the CPU and DASD capacity of the system, but, generally speaking, less than 100 users and less than 600K records (potentially spanned before V6.3) should not represent a problem for any given customer.

If the number of users and total records seem reasonable, then the collection of some performance figures is mandatory. Measurements may be taken by using RMF or similar tools. You may want to use TASID, which is an easy-to-use tool available from MVSTOOLS. TASID provides snapshots of CPU status, LPAR usage, number of TSO users and VSAM overall performance figures.

Sometimes the solution to the problem may need a combined effort of MVS system programmers and Information/Management specialists. If you think that there are strong possibilities of this happening, make sure that MVS system programmers also collect performance data of the system at the same time intervals you collect data for Information/Management

1.3 Latest Performance Enhancements

The most significant performance enhancement in Version 1.1 is the availability of multiple SDIDS support. In Version 1.1 the SDIDS can be split into multiple VSAM clusters. This option can help reduce contention in the SDIDS. In order to activate the function, the administrator should perform an analysis of the key distribution in the SDIDS and split the data set using key ranges.

Maybe the most significant performance enhancement in Version 6.3 is related with the option to eliminate spanned SDIDS records. Nevertheless, it is worth highlighting all performance improvements, which are:

- Elimination of VSAM spanned SDIDS records
- SDDS storage capacity

You can find below a more detailed explanation for each of the Version 6.3 performance enhancements.

Elimination of VSAM spanned SDIDS records

Significant enhancements have been made to eliminate the need for VSAM spanned records for the SDIDS. VSAM spanning is eliminated by splitting a single logical SDIDS record into multiple physical records for throughput improvement. You can, if there is a need for it, continue to use spanned records. The available key lengths for the SDIDS are 16, 18, 32 or 34 bytes.

SDDS storage capacity

The maximum storage capacity of SDDS is extended to a maximum of 100 VSAM clusters (400 GB), which allows for more interleaving of ENQueues for SDDS access.

Multisystem Database Access Enhancements

A timer function is implemented to terminate waiting after 30 seconds, allowing the process to continue when Multisystem Database Access (MSDA) is hung while waiting for a response from APPC/MVS.

An error message is written out and the APPC conversation is reallocated following a timeout failure. This brings up a performance improvement and a simplified problem determination.

Queued Notification Mail Processing

Notification mail can be queued to the BLX-SP as an alternative to processing mail immediately in the user′s address space, thereby improving the performance by reducing processing at record file time.

BLGUT3 and BLGUT4 utilities
BLGUT3 and BLGUT4 database utilities now offer improved performance. Changes have been made to prune multiple updates for a record in the offloaded log data set, so that records will only be processed twice, no matter how many updates were made.

1.4 Tuning Information/Management

Tuning is a complex area. Sometimes it can have a see-saw effect in that changing one parameter will achieve what you set out to do, but at the cost of creating a bottleneck elsewhere.

Tip

One golden rule of performance management is to change just one thing at a time, then to measure the effects of that change before proceeding to make further changes.

Since Version 5.1, Information/Management has utilized a single address space to manage all I/O activity, bringing the benefits of:

- Increased throughput
- Improved data integrity
- Improved response times

Note: RMF can show user activity by means of figures from I/O measurements.

This redbook discusses the following topics in the context of performance:

1. Information/Management Parameters - This chapter looks at the most important Information/Management parameters that can influence performance. Some of these are only available since &Infover..
2. I/O Factors - This chapter discusses various aspects of VSAM, placement of data sets, and how to get the best out of it.
3. Memory - This chapter looks at memory considerations.
4. CPU - This chapter looks at CPU considerations.
5. Version 6 Considerations - This chapter looks at how the new features that are available with Version 6, and how they can affect your performance.
6. Customer Study

Chapter 2. Tuning Information/Management Parameters

This section describes parameters that can be coded in the session-parameter member to influence the performance of Information/Management. The following parameters are discussed:

SRCHLIMIT (SRCHENQ) Coded in BLGCLUST macro

COGENQ Coded in BLGCLUST macro

SORTPFX Coded in BLGPARMS macro

PNLBCNT Coded in BLGPARMS macro

If you have made any changes to your session-parameter member, then it should be assembled into a library that is either in the MVS linklist, or accessible to the users via ISPLLIB or STEPLIB. A sample job, BLGALSPM, is provided in SBLMSAMP for this.

Please remember when you assemble a session-parameter member that it is a nonreentrant module, and therefore you should include NORENT as one of your linkage editor parameters.

For the following discussion, you should consider a system where many Information/Management users are working concurrently, each one contributing transactions to the overall system load. As each transaction is processed, it will access the SDIDS. This access to the SDIDS will be serialized through the use of enqueues.

Write transactions have exclusive access to SDIDS whereas read transactions allow concurrent access. The number of write operations is controlled by the parameter COGENQ in the session parameter. The number of read operations is controlled by the first parameter of the search limit parameter.

Some transactions will keep the enqueue briefly, but others might need a longer time to complete. Older Information/Management versions (V5 without APAR OY62656, Version 4 and lower) kept the enqueue for the entire transaction elapsed time, inhibiting competing transactions to execute while waiting for the long-running transaction to finish and reducing overall system throughput. Information/Management, since Version 5.1, provides mechanisms to handle more enqueue requests, allowing for a better sharing of the SDIDS and improving the overall system throughput.

Tip

Every Information/Management user reads a session-parameter member at session startup. You can create multiple, unique session-parameter members and make them available in your system. You can group users according to their activity, based on transaction type, and set up session-parameter members that are appropriate for each group. You can provide different Information/Management startup CLISTs (or REXX EXECs), automating the selection process.

2.1 Tuning SRCHLIMIT ENQUE (or SRCHENQ)

The SRCHENQ parameter from earlier versions of Information/Management was expanded in Version 6 by means of the SRCHLIMIT parameter, which is based on SDIDS counts. It allows you to fine-tune the interleaving of *range read enqueues*. (For an explanation of this concept please refer to A.1.3, "What Can Cause Massive Numbers of SDIDS Reads?" on page 82.)

Read operations during one transaction are counted separately. For example, if you set up a search argument such as:

```
SE STAC|O. PH|.
```

and you have 50 different values of STAC|O., 90 different values of PH|., and have set your first parameter of the search limit to 100 the read operation is *not* interrupted.

SRCHLIMIT is a three-value parameter. The first value works like SRCHENQ, the second value allows you to send a warning to any user that exceeds a specified number of SDIDS range reads, and the third value allows you to set a maximum number of SDIDS range reads for search or sort operations. You can use these new features provided with the SRCHLIMIT parameter to identify transactions that use system resources inefficiently.

When tuning SRCHLIMIT ENQUE, the first value of the SRCHLIMIT parameter, keep the following in mind:

1. Never set to a number less than 10. As continuing the transaction after each enqueue operation requires the position in the SDIDS to be re-established, thereby increasing the I/O to the file. Setting SRCHLIMIT ENQUE to 10 adds one read operation for each 10 originally required. (That is a 10% increase in the number of I/O operations required to complete the transaction.)
2. Setting to 100 will only add 1% to the current transaction I/O count, and provide frequent opportunities for other transactions to get access to the SDIDS. This will tremendously improve response times, especially for transactions that require only a few enqueue opportunities to complete.
3. There is CPU usage performing each ENQ/DEQ operation.

Please note that the SRCHENQ parameter has not been removed, but the SRCHLIMIT parameter will take precedence over SRCHENQ if they are both present, but with different values. The SRCHENQ parameter has remained for compatibility purposes with earlier releases of Information/Management.

See Appendix A, "Searching and Sorting Flash" on page 81 for a further discussion of searching and sorting, and how to avoid transaction performance degradation when performing those operations.

Tip

Avoid big-range searches, as these cause a lot of I/O to occur. (To learn how to determine if you have this problem, please refer to A.1.5, "What Can I Do?" on page 84.)

2.2 Tuning COGENQ

This parameter was added to Version 5.1 by the fix to APAR OY62656. Since Version 6.1 is part of the base product. It will help the performance of file transactions at a busy time, particularly if Multisystem Database Access has been implemented.

Add this parameter to your session-parameter member to specify the number of fields to be cognized during a single enqueue when filing a record.

Information/Management Version 5.1 obtains and releases an exclusive enqueue for each field being cognized. This creates many enqueue opportunities for competing transactions.

COGENQ allows you to hold the enqueue while multiple fields are being cognized, which reduces the total number of enqueues, but makes each individual enqueue last longer. The COGENQ value specifies how many fields are to be cognized while you hold the enqueue.

> **Tip**
>
> COGENQ allows you to bring in a degree of prioritization, so that an interactive user could have a different session-parameter member than an API (batch) user. Assigning the interactive user a session-parameter member with a higher value for COGENQ increases his or her priority. A help desk function that generates a lot of create activity, but not much update activity, could benefit from a high COGENQ value.

There are several choices in the way that COGENQ is set. These are as follows:

- Set COGENQ to zero, or a very large number (probably greater than 50).

  This will cause the enqueue to be held for the duration of all cognize and uncognize activity. This gives all power to file activity, and others will be locked out at busy times. MSDA users should use a value of 0, to allow one ENQ/DEQ instead of multiple ENQ/DEQ operations to save a record which would cause complex wide traffic for the multiple ENQ/DEQs necessary to file records.

- Set COGENQ to 1.

  The ENQ will be released between the cognize of each field. This is appropriate for batch users, who won′t care if they have to give way to other users, as their response time is less important.

- Set COGENQ to a number between 1 and 50.

  The priority of the file activity is proportional to the value of COGENQ. So COGENQ=20 gives a higher priority than COGENQ=10 since the ENQ is held twice as long.

On light loads, COGENQ will have little effect. It is at busier times that the benefits of managing the enqueue contention will be seen, provided that it is set to a value greater than 1.

The diagram in Figure 1 on page 8 shows how COGENQ will influence both those using the file, and those doing a range search of 3000 records. With a high value for COGENQ, file activity is much quicker at the expense of the user running the search.

In this chart, when COGENQ is 100, this gives the file user a response time of about a second, whereas the range read has a response time of 3.5 seconds. When COGENQ is set to 1, then this is a no-win situation, as both types of users have a response time of about 2.7 seconds. The optimal value in this example is to set COGENQ to 10, which gives the file user performance almost as good as having COGENQ=100, and at the same time does not impact the range-read user too greatly.

Note: COGENQ=0 means no enqueue release during the file operation. It shows in the chart as the topmost COGENQ point in the curve.

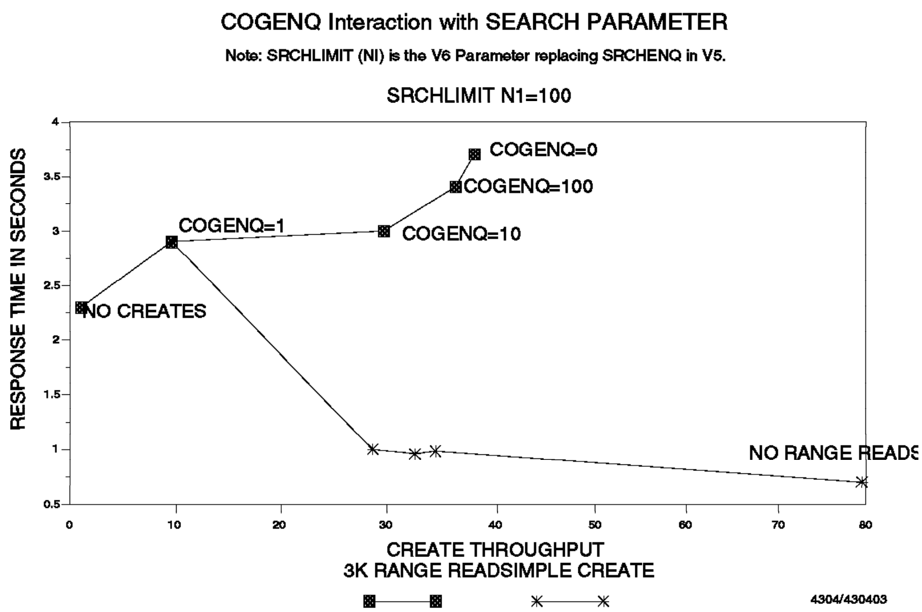

Figure 1. How COGENQ Interacts with SRCHLIMIT N1

Tip

This value must be specified according to the priority we want to give to the user activities. If the user creates records most of the time (for example, help desk personnel) then COGENQ should be set to zero or a value higher than 10 in the user session-parameter member. The default value (10) is good for balancing the search and file activities. It should be used for general users. A smaller value (1) should be used in the session-parameter member assigned to batch processes (such as an API program loading records to Information/Management continuously).

2.3 Tuning SORTPFX (Related to Search Result List)

When you want to tune search result list (SRL) output, you had better take into account the following three key SORTPFX parameters:

- N1 - This value determines the maximum number of matches the search result list could contain. The default is 2 147 483 647.
- N2 - This is the maximum number of matches in the search results list that will be sorted with the sort-on-prefix function. The default is 32767.
- N3 - This is the maximum number of lines in a search results list (SRL) that can be sorted using an effective internal sort routine. The default is 500.

Tip

Setting N1 to 1000 (max) gives a maximum of 1000 matches. Setting N2 to 100 gives a sorted output if your search argument results in less than 100 matches (which is four screens). Any results between 100 and 1000 gives an unsorted SRL.

The above parameters are compared to the length of your SRL. If your SRL exceeds the N1 value, then no records are displayed but the result (number of hits) of the search is displayed as a message to the user. However, if the number of matches is less than N1 but greater than 32767, then Information/Management will display 32767 matches as this is the highest number that N2 can be set to. Please refer to *Planning and Installation Guide and Reference*, SC34-4603, for a discussion of the SORTPFX affect the sorting of the SRL in the BLGPARMS macro.

The two sort methodologies are:

- If the number of matches in the SRL is less than N2, and the sort-on-prefix is specified on the table panel, then Information/Management will sort the SRL before displaying the SRL. This value is ignored if the sort-on-prefix is not specified on the table panel used for the display. If the number of matches exceeds the N2 value then no sort-on-prefix is done even if the table panel that displays the data may require it.
- If the SRL length is greater than N3, then the SDDS is not read in the sorting process. Instead, all the SDIDS records related to the search are read. Since the records are already sorted, they can be used to retrieve the SDDS records in the proper order for the first viewable screen, etc. That′s why it′s very important not to sort on fields (such as RNID/) that can have thousands of SDIDS records.

Tip

The value of N2 should be the same as N3, where there is a possibility of the SRL being very large. This is because large range internal sorts, specified by N3, will diminish the speed advantage offered by the internal sort. The values should be set as low as possible.

2.4 Tuning PNLBCNT

This is the panel buffer count parameter. It allows you to control the number of local panel buffers that are available to each user for storing the most frequently used panels. A least recently used algorithm is used to discard the panels that have remained unreferenced for the longest period of time.

This parameter has a default value of 50. However, it can be changed to any positive integer between 1 and 32767.

Information/Management does not allocate buffer storage until it′s needed, so you can specify 32767 and Information/Management will only allocate storage for the number of panels actually used during the session. The BLG22556 message tells you how many were actually allocated.

To tune this requires analysis of the BLX-SP log for each user. When a user quits Information/Management, a BLG22556 message is written to the log detailing how many panels were picked up from these buffers without needing BLX-SP to initiate I/O. As of Version 6.1 log information gives the number of buffers allocated for each user and the maximum used, so this information can be used to determine if PNLBCNT should be raised or lowered.

Tip

We recommend that you start with the default value for PNLBCNT, and tune it later aiming for between 50% and 75% reads to be satisfied from the buffers, although there will be a storage cost to achieve this.

Message panels, TSPs, and in some special cases control panels must be read from the panel data set for each use. If you′re using a large number of TSPs, a somewhat lower percentage of panel reads will be satisfied from local panel buffers.

Chapter 3. I/O Performance Tuning and Capacity Assessment

It is important to look at I/O, as this is usually a bigger performance constraint than CPU. This section describes actions that can be taken to make the I/O subsystem work efficiently, and in turn optimize your Information/Management performance.

You can find some performance figures in chapter 4 of the *Planning and Installation Guide and Reference*, SC34-4603 manual.

3.1 VSAM Data Sets

Information/Management uses seven different types of data sets:

- VSAM data sets:
  - The Structured Description Data Set (SDDS)
  - The Structured Description Index Data Set (SDIDS)
  - The Structured Description Log Data Set (SDLDS)
  - The Dictionary Data Set (DICTDS)
  - The Read Panel Data Set (RPANLDS)
  - The Write Panel Data Set (WPANLDS)
- Partitioned data sets:
  - The Report Format Data Set (RFTDS)

Out of the above, the SDDS, the SDIDS and the SDLDS data sets are the most important to take into account when looking at performance issues. However, there are also some considerations and recommendations to keep in mind for the remaining data sets.

3.1.1 SDIDS Spanned Records No Longer an Issue

You should consider using the 18 or 34-byte SDIDS key length available since Version 6.3 of Information/Management to eliminate the need to span VSAM records and to significantly improve throughput and performance.

Figures taken from actual customer installations show a performance gain of 30%.

3.1.2 SDDS Considerations

The SDDS is a VSAM key sequenced data set. It is the data set that contains the data you enter into an Information/Management database.

Since Version 6.3 the SDDS can have from 1 up to 100 VSAM key-sequenced data sets (clusters) that contain data records. A 100-clusters for the SDDS would allow 400 GB of information in the database.

The data collected in the SDDS includes not only the field contents that you enter, but other data too (for example, who put it in, when, and which panels they were on when they entered the data). Field content is also linked to the field it was meant for. So, the data actually defines not only the contents but the structure of the data model for the database. This data model of the database is dynamic, so you can use PMF to add or delete fields to your panels and Information/Management can still figure out what is in the SDDS.

The SDDS may be structured differently depending on whether you choose to use a 7- or 8-byte key. Very large SDDS records may be broken into multiple physical VSAM records if necessary (although Information/Management will treat them as one logical record). The keys of the SDDS records are numbers, and the lowest available number is always assigned to each new record.

The following are the factors that limit the number of SDDS records in Version 6.3, Version 6.2, Version 6.1 as well as in Version 5.1:

- The Information/Management database size per volume.

  The database size is limited by VSAM. The maximum size of a VSAM data set is 4 GB. All previous releases to Information/Management Version 5.1 were limited to a maximum database size of 4 GB. Version 5.1, Version 6.1 and Version 6.2 can have up to five clusters viewed as a single logical SDDS, which allow for SDDS sizes of up to 20 GB, whereas with Version 6.3 a 100-cluster SDDS would allow up to 400 GB.

- The record size.

  Assuming that your SDDS average record size is 4 KB (large lists or long freeform text data will generate even larger records) the number of records it can hold will be 1.3 million if using a 3380 disk. With an average record size of 2 KB the number will duplicate, about 2.6 million. Again, these figures depend on your average record size.

  If you′re using a five-cluster SDDS, you can, in theory, use five full volumes. Assuming they hold two million records each, it makes a total of ten million records. For a 100-cluster SDDS the theoretical figure skyrocket up to two hundred million records.

- The SDIDS maximum record size.

  If you are using 16 or 32-byte key lengths, the maximum record size for the SDIDS determines how many SDDS record references an SDIDS Master Bit list can have.

  Note: The current maximum limit of logical SDDS records for a spanned SDIDS with a 16 or 32 byte key length is 16 777 215.

  If you are using 18- or 34-byte key lengths, the maximum record size of the SDIDS no longer limits the number of records that an SDDS may have from a practical point of view. The minimum number of SDDS records supported is over 125 million when using a SDIDS maximum record size of 507 bytes. The maximum is over 8.6 billion SDDS records when using an SDIDS maximum record size of 32 760. Generally speaking, you will run out of physical space to store the SDDS records before you reach any limit imposed by the maximum record size of the SDIDS when using the 18- or 34-byte SDIDS key length.

For planning purposes, once you have estimated the maximum number of SDDS records to be kept, the rule of thumb is:

- If the estimate is below 261 000, then using 18 or 34-byte SDIDS key length is not a major issue. (16 and 32-byte will also work.)

- If the estimate is equal to or over 261 000, then using 18 or 34-byte SDIDS key length is highly recommended. The 16 or 32-byte key lengths are not recommended in this situation.

Although you may always want to use 18 or 34-byte key lengths, you can use the following formula to compute the maximum record length for the SDIDS when you need to deal with 16 or 32-byte key lengths:

```
SDIDS maxlrecl = (Number of SDDS records/8) + SDIDS key length + 8
```

To know the maximum number of SDDS records supported by a given SDIDS maxlrecl, you may want to use one of the following formulas:

- For 16 or 32-byte SDIDS key lengths:

  ```
  Number of SDDS records = (SDIDS maxlrecl - 8 - SDIDS keylen) * 8
  ```

- For 18 or 34-byte SDIDS key lengths:

  ```
  Number of SDDS records = ((SDIDS maxlrecl - 8 - SDIDS keylen) * 8) * 32767
  ```

  Where the maxlrecl value ranges between 507 to 32760, use a maxlrecl which is seven less than the CI size, ensuring that the CI size is optimum (bytes per track) for the target storage device.

Planning and Installation Guide and Reference, SC34-4603.

3.1.3 SDIDS Considerations

Important

Since Version 6.3 Information/Management allows for the elimination of VSAM spanned records for the SDIDS. In addition, Version 6.3 will no longer limit the number of database records because of the SDIDS maximum record size.

The SDIDS is a VSAM key-sequenced data set that consists of a single VSAM cluster. You only define one SDIDS for the SDDS, regardless of whether the SDDS is 1 or 100 VSAM clusters.

The SDIDS is a special high-performance index that accelerates the search for records in the SDDS. The SDIDS is sometimes called the Glossary. The SDIDS can be viewed as being a collection of records that represent every searchable condition in the database. The SDIDS record for each searchable entity contains a bit string where the relative location of all the 1 bits equals the keys to all SDDS records that match the searchable condition. For example, the SDIDS record keyed PERS/JONES represents a searchable entity that stands for all records where the submitter is JONES.

When a record is filed, the amount of I/O against the SDIDS will be influenced by the number of fields that are cognized. For every field that is fully cognized, there might be up to four I/O operations (one read, one write for the p-word and one read, as well as one write for the s-word) to the SDIDS. By using p-word cognize only you can divide by two the amount of I/Os per field (one read and one write for the s-word will be avoided).

You can decide what fields are cognizable, thus reducing the I/O by reducing the number of fields to be cognized for each SDDS record filed.

```
  Environment   Dialog   Record   Window   Options   TSO   Help
 ------------------------------------------------------------------- MORE: +-
 ===>

 BLG1TGLD                         GLOSSARY DISPLAY

  DATABASE: 5

                       COUNT  KEYWORD

                           1 TIME/13:36
                           1 TIME/13:38
                           1 TIME/13:53
                           1 TIME/13:56
                           1 TIME/13:58
                           1 TIME/14:20
                           1 TIME/14:29
                           1 TIME/14:32
                           1 TIME/14:34
                           1 TIME/14:35
                           1 TIME/15:17

  LINE COMMANDS:  NONE
  SCROLL, OR END TO TERMINATE
                                                                        IDLE
```

Figure 2. Information/Management Glossary Display Panel

If you look at the glossary, Option 8 on the main panel, you can get an idea of what the SDIDS contains (see Figure 2). Each line in the glossary is a valid key value for an SDIDS record. The number on each line represents the count of bits set to one (1) in that SDIDS record. (You cannot see directly the bit string in each SDIDS record because it is compressed.) The relative offset of each one bit in the bit string represents the relative record number (key) of an SDDS record that contains the item related to that SDIDS key value.

Normally, there are only a few Information/Management logical records referenced by any SDIDS record (search word). This means that most SDIDS records are compressed to a relatively short length. Although some records cannot be compressed and are very large, the average record size can be as small as 50 bytes for a 16 or 18-byte key or 65 bytes for a 32-byte key. The average record size is dependent on the contents of the database, so most SDIDSs will have an average size that is larger than 50 or 65 bytes. If you have a previous version and have records established, you can run BLGUT21 to get your average record length.

The maximum record size of the SDIDS determines the length of the Master Bit List. If you are not using spanned records, the maximum record size is 32 760 with a CI size of 32 768 regardless of the key length specified (16, 18, 32 or 34). *Planning and Installation Guide and Reference*, SC34-4603.

3.1.3.1 Dealing with MAXRECSIZE for SDIDS

When using an SDIDS with 18 or 34-byte key lengths, the MAXRECSIZE of the SDIDS no longer limits the number of records in the SDDS.

Factors to be considered when assessing the right MAXRECSIZE for the SDIDS, when either 18 or 34 byte key lengths are used, include:

- Optimum DASD utilization

The goal is to minimize the amount of unusable space per track, which depends on the DASD track capacity. Performance will also benefit from it.

- LSR storage buffer capacity
- VSAM I/O activity

  Smaller CISIZE may require more I/O requests for database search activity, thus affecting overall performance. Database update activity processes SDIDS data at the segment level, thus shorter CISIZE might yield better performance.

- SDIDS storage capacity

  The SDIDS can only hold 4 GB of data, each physical SDIDS record has a fixed amount of overhead associated with it, therefore a short MAXRECSIZE would require a larger amount of overhead than a large MAXRECSIZE. This is only a factor if there is a high number of cognized records in the SDIDS.

With the 18 byte key SDIDS architecture, the maximum possible number of SDDS records that can be represented is:

```
MAX_SDDS_RECORDS = 32734 (max number of bytes for a bitlist) x
  32767 (max number of physical segments for a logical SDIDS record) x
  8 (number of bits in a byte)
  = 8 580 759 840 (or, approximately 8.6 billion records)
```

This formula takes into account that the maximum SDIDS record length that Information/Management will use is 32760 which would require a CISIZE of 32767. The number 32734 is 32760-26. (Each SDIDS record contains 26 bytes of overhead data, including 18 bytes for the key.)

As you can see, the SDIDS has the capacity to accommodate 8.6 billion SDDS records, 100 SDDS clusters do not provide sufficient storage capacity to contain that number of records.

From a pure Information/Management standpoint, the preferred SDIDS CISIZE would be 32,767. Though that answer doesn't take DASD track capacity into consideration, it does yield the fewest physical SDIDS records, the master bit list would require seven segments. However, it is highly recommended to reduce this number to achieve optimum DASD utilization.

3.1.3.2 Recommended SDIDS Key Lengths

Information/Management supports four sizes of SDIDS key lengths: 16, 18, 32 and 34 bytes. The 18 and 34 bytes are provided since Version 6.3 to improve database performance and to eliminate the need for VSAM spanned records.

Important

It is highly recommended to use the 18 or 34-byte SDIDS keys to eliminate the need for spanned SDIDS records. By doing this you will significantly improve throughput and performance.

You specify the key when defining the SDIDS. The key size is detected when Information/Management starts. You must run the BLGUT1 utility to restore SDIDS data sets when you want to use a new key size.

3.1.4 SDLDS Considerations

This data set contains copies of your SDDS records that have been entered, updated, or deleted since the SDLDS was last offloaded. The records are not compressed, and look almost exactly like the original SDDS records. One full copy of the record will be stored in the log for each file (or delete) operation.

The size of this data set is a consideration, and will be affected by the number of changes made to the database each day. Your log data set needs to be large enough to hold one copy of any record filed or deleted in between runs of BLGUT4.

If you use Automatic Log Save, available in versions higher than Version 6.1, then this may allow you to reduce the size of the SDLDS, as the offload process will be more frequent.

You can also operate without this data set which would improve file time response, but then forward recovery could be carried out should the need arrise.

3.1.5 DICTDS Considerations

The dictionary data set (DICTDS) contains the structured words (s-words) and prefix words (p-words) that are used to search the SDIDS and to control the display data on Information/Management panels for all databases. It also contains validation patterns that are used to control the syntax of values entered into Information/Management panels.

The DICTDS can use Local Shared Resources (LSR) buffers. However, if there are many reports running during an online period, then the LSR buffers allocated to the DICTDS should not be shared with other Information/Management data sets. This is because a heavy I/O load can build up when the LSR buffers are being used by the other data set(s) in the same LSR pool.

3.1.6 Panel Data Set Considerations

These are usually used in read mode, and are unlikely to affect performance if LSR is used, as there are no exclusive ENQs involved, so several users can be using the same panels simultaneously without impacting each other (assuming no other DASD restrictions are involved).

Combining panels into one data set improved response as this saves time opening, searching and closing many panel data sets to retrieve necessary panels.

Information/Management has a set of user panel buffers, reserved in each user address space for frequently used panels, to allow for those to be reused without extra I/O. The parameter that controls this is PNLBCNT. This parameter is set in the session-parameter member, and the default value is 50, meaning that 50 panels can be kept in the buffers. Please refer to 2.4, “Tuning PNLBCNT” on page 10 for how to tune this parameter.

Tips

- Place the SDIDS index component on a cached volume with DASD FastWrite.
- Place the SDDS index component on a cached volume with DASD FastWrite.
- For recovery and performance, place the SDDS and SDIDS on different volumes.
- Use LSR to keep the entire index component of the read panels data set, SDIDS, SDDS, and part of the data portions, in storage.
- If possible use LSR for the SDIDS data portion as well.
- Don′t cache the low-use data sets, as this will waste a valuable resource.
- Reorganize a customer-modified read panels data set before placing it into production if there are CI or CA splits due to modification or updating of panels.

3.2 Optimizing Control Interval Sizes

When the value for the SDDS CISIZE is chosen, the purpose is to set this value such that the longest physical record will fit into one control interval. Also, we want to fit two smaller logical records into a shared control interval if possible, so that we at best eliminate CI splits, and at worst to keep them to a minimum.

To optimize the values for the CI sizes for the three important data sets, the SDDS, the SDIDS, and the panels data sets, proceed as follows:

1. Run BLGUT20 to obtain statistics about the SDDS records. You can see an example in Figure 3 on page 18.

   The output of BLGUT20 gives you a discrete distribution of the SDDS record sizes. As in any other statistical report, you must use your own experience and judgement before making up your mind. However, it is always very useful to find out more about the distribution shape and characteristics. You can find below some guidance and advice:

   - Check whether the average SDDS logical record size is closer to the medium. The report doesn′t give you the medium, but gives you numbers depending on CI sizes.

     In the example in Figure 3 on page 18 you can see that the average record size is 2174 and that most of the SDDS records require a 2048 or 3072 CI size. That means that the medium is closer to the average.

   - Check whether the distribution gives you enough data to figure out the distribution shape. This might be the case if you have many records requiring a greater than 8192 CI size.

     Should the above be your case, you have to find out something more about the distribution or sizes above 8192 bytes. A good starting point is by assuming sizes for records in the different CI sizes intervals, and, from there, calculating and comparing the resulting average record size against the average SDDS logical record size in the report. If the results don′t match very well, you can keep making assumptions with regard to the average record size over 8192 sizes by taking into account the maximum SDDS logical record size from the report.

```
 . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .
 SDSF OUTPUT DISPLAY INFO4I   JOB05963  DSID   101 LINE 1      COLUMNS 01- 80
 COMMAND INPUT ===>                                           SCROLL ===> CSR
1SDDS DATA SET ANALYZED:  'INFOMAN.MIG42.SDDS'     TIME: 23:34:14 DATE: 11/24/94
 CI SIZE: 4096     KEY LENGTH: 7     MAXIMUM RECORD LENGTH: 4089
0TOTAL DATA SET SIZE: 100022
     TOTAL SDDS STRUCTURED DATA SIZE: 74894
     TOTAL SDDS STRUCTURED TEXT SIZE: 10060
     TOTAL SDDS STRUCTURED HISTORY SIZE: 15068
 LOGICAL RECORDS PROCESSED: 46
 PHYSICAL RECORDS PROCESSED: 48
0MINIMUM SDDS LOGICAL RECORD SIZE: 1013
 AVERAGE SDDS LOGICAL RECORD SIZE: 2174
 MAXIMUM SDDS LOGICAL RECORD SIZE: 10736
0SDDS RECORDS REQUIRING A  512 CI SIZE: 0
 SDDS RECORDS REQUIRING A 1024 CI SIZE: 1
 SDDS RECORDS REQUIRING A 2048 CI SIZE: 30
 SDDS RECORDS REQUIRING A 3072 CI SIZE: 14
 SDDS RECORDS REQUIRING A 4096 CI SIZE: 0
 SDDS RECORDS REQUIRING A 6144 CI SIZE: 0
 SDDS RECORDS REQUIRING A 8192 CI SIZE: 0
 SDDS RECORDS REQUIRING A GREATER THAN 8192 CI SIZE: 1
 SDDS RECORDS REQUIRING MULTIPLE CONTROL INTERVALS: 1
0SDDS ANALYZE UTILITY COMPLETE
```

Figure 3. Information/Management BLGUT20 Utility Program Output

As an overall rule of thumb, set your CISIZE to two or three times the average value according to these statistics.

2. Run BLGUT21 to obtain statistics about the number of records in a control interval for the SDIDS. Choose to set your CISIZE to 20% more than the largest value according to these statistics (refer to Figure 4).

```
SDIDS DATA SET ANALYZED:  'INFOMAN.MIG42.SDIDS'     TIME: 23:44:31 DATE: 11/24/9
CI SIZE: 2048     KEY LENGTH: 32     MAXIMUM RECORD LENGTH: 2041
CURRENT NUMBER OF DATABASE RECORDS: 46
MAXIMUM NUMBER OF DATABASE RECORDS ALLOWED: 16008
TOTAL DATA SET SIZE: 15258
SDIDS RECORDS PROCESSED: 321
MINIMUM SDIDS RECORD SIZE: 41
AVERAGE SDIDS RECORD SIZE: 48
MAXIMUM SDIDS RECORD SIZE: 150
SDIDS SPANNED RECORDS REQUIRING MULTIPLE CONTROL INTERVALS: 0
SDIDS ANALYZE UTILITY COMPLETE
```

Figure 4. Information/Management BLGUT21 Utility Program Output

Tip

Increase the CI-size for the SDIDS up to 32 KB, according to the needs of your installation. If you have to go spanned, it is useful to decrease the CI-size to 26 KB because that represents half a track of a 3390 disk, and all DASD/MVS algorithms are tuned for this size. Again, spanned records are not an issue in Version 6.3, but, if you need them, you better consider setting CI-size to 26 KB, as this gives better performance, although generates more spanned records. Remember to use a cisize that yields the best DASD utilization for the device type you are using for the Index.

3. Run BLGUT22 to obtain statistics about the number of records in a control interval for the panel data sets. Choose to set your CISIZE to the largest value according to these statistics, rounded up to the nearest 2 KB (refer to Figure 5 on page 19).

   Note: Be aware that there is no further distribution intervals for records larger than 8 KB. This applies to all clusters, as well as BLGUT20, BLGUT21 and BLGUT22 utilities.

```
PANEL DATA SET ANALYZED:  'INFOMAN.MIG42.WPANELS'     TIME: 00:05:01 DATE: 11/25
CI SIZE: 4096     KEY LENGTH: 10     MAXIMUM RECORD LENGTH: 2041
TOTAL DATA SET SIZE: 6417
LOGICAL RECORDS PROCESSED: 4
PHYSICAL RECORDS PROCESSED: 5
MINIMUM PANEL SIZE: 482
AVERAGE PANEL SIZE: 1604
MAXIMUM PANEL SIZE: 4030
PANELS REQUIRING A  512 CI SIZE: 1
PANELS REQUIRING A 1024 CI SIZE: 1
PANELS REQUIRING A 2048 CI SIZE: 1
PANELS REQUIRING A 3072 CI SIZE: 0
PANELS REQUIRING A 4096 CI SIZE: 1
PANELS REQUIRING A 6144 CI SIZE: 0
PANELS REQUIRING A 8192 CI SIZE: 0
PANELS REQUIRING A GREATER THAN 8192 CI SIZE: 0
PANELS REQUIRING MULTIPLE CONTROL INTERVALS: 0
PANEL ANALYZE UTILITY COMPLETE
```

Figure 5. Information/Management BLGUT22 Utility Program Output

4. Monitor the CI splits using the above listed utilities or the Access Method Services IDCAMS LISTCAT.

   Note: LISTCAT will only give accurate results if the VSAM data set is closed, when Information/Management has been stopped.

   The representative workload we defined contains no BATCH, API, sorting, or other complex automation (TSPs, user exits, etc.) that some customers do use. CPU costs for some transactions not included in the workload can be substantially higher in Version 5.1 than in Version 4.

Tip

CI/CA splits heavily decrease your VSAM performance. The only chance in a spanned SDIDS environment is to run delete/define SDIDS cluster and run BLGUT1 afterwards.

In all cases it is highly recommended to define CI sizes that will be the optimum considering the following criteria:

- Best size in terms on Information/Management logical records
- Best fit, according to the physical characteristics of the DASD device

You can find in Table 1 on page 20 the different bytes per track of three types of disk storage devices.

Table 1. Bytes per Track for Various DASD Types	
Device Type	**Bytes per track**
3380	47 476
3390	56 664
RAMAC	56 664

Generally speaking, even though CI sizes of 32 768 may be desirable from an Information/Management point, it is wise to optimize CI sizes based upon disk characteristics.

You can find in Table 2 some figures related to CI sizes and disk space utilization. The figures are based on the algorithms used for VSAM space utilization and optimization.

Table 2 (Page 1 of 2). Assessment of CI Sizes and Related Disk Space Utilization

CI sizes	**Block Size for 3380** **Space Utilization**	**Block Size for 3390** **Space Utilization**	**Block Size for RAMAC** **Space Utilization**
32 768	8 192 86% usage	10 240 91% usage	10 240 91% usage
30 720	6 144 91% usage	10 240 91% usage	10 240 91% usage
28 672	14 336 91% usage	7 168 89% usage	7 168 89% usage
26 624	6 656 84% usage	26 624 94% usage	26 624 94% usage
24 576	6 144 91% usage	24 576 87% usage	24 576 87% usage
22 528	22 528 95% usage	5 632 90% usage	5 632 90% usage
20 480	20 480 87% usage	10 240 91% usage	10 240 91% usage
18 432	6 144 91% usage	18 432 98% usage	18 432 98% usage
16 384	8 192 86% usage	16 384 87% usage	16 384 87% usage
14 334	14 334 91% usage	7 168 89% usage	7 168 89% usage
12 288	6 144 91% usage	12 288 87% usage	12 288 87% usage
8 192	8 192 86% usage	8 192 87% usage	8 192 87% usage
4 096	4 096 86% usage	4 096 87% usage	4 096 87% usage

Table 2 (Page 2 of 2). Assessment of CI Sizes and Related Disk Space Utilization			
CI sizes	**Block Size for 3380** **Space Utilization**	**Block Size for 3390** **Space Utilization**	**Block Size for RAMAC** **Space Utilization**
2048	2048 78% usage	2048 76% usage	2048 76% usage
1024	1024 67% usage	1024 60% usage	1024 60% usage
512	512 50% usage	512 45% usage	512 45% usage

Tip

The best performance is reached when optimum CI sizes are used. For instance, optimum performance for 3390 and RAMAC can be achieved when using CI sizes of:

- 26624 (95%)
- 18432 (98%)

3.3 LSR Buffers

This is a very important area for tuning Information/Management, as there is a lot of file activity involved. The aim is to make enough buffers available to avoid a large percentage of physical reads to the data set. There is a cost to having these buffers, and that cost is in the virtual storage overhead.

When Information/Management is used without the Multisystem Database Access feature, it works like Version 5.1. In this situation no buffer invalidation needs to occur if an update, create, or delete is done, which is good from a performance point of view. This clearly is of particular benefit to the SDDS and SDIDS, since they are the most highly used data sets.

Particular attention should be paid to the values for the buffers associated with the index components of the data sets. A buffer hit rate of 100% should be achievable at no great cost as the records are small, but the gain is large as they are highly used.

The data component of the SDIDS is a good candidate to be highly buffered for the following reasons:

1. The records are small.
2. The SDIDS data component is the highest contributor to enqueue delays.
3. The records are frequently accessed.

The SDDS will benefit less from the buffering, for the reasons below, but it will still be worthwhile to have enough to handle a few hundred control intervals.

1. The records are large.
2. The records are less frequently accessed.
3. Random access.

Tip

The higher the create/update activity, the lower your LSR buffer hit ratio will be.

3.4 VSAM Buffers Checklist

The following checklist should help to optimize the use of LSR buffers:

1. Ensure that the optimum CI size is defined for your VSAM files. This should be done as explained in 3.2, "Optimizing Control Interval Sizes" on page 17.
2. If you increase the CI size of any data or index component, you must review your LSR definitions, change them, and reassemble the BLXVDEF.
3. Ensure that your LSR definition is appropriate when changing the key lengths of the SDDS or SDIDS data sets.
4. Give your dictionary and panel data sets the same CI size for their data components, to allow them to share buffers.
5. Aim for a 100% hit rate on your index components, as this will keep the information permanently in storage.
6. Avoid NSR except for the SDLDS.
7. Don′t define too many NSR buffers or set up STRNO too high; it will waste valuable CSA below the 16 MB line.
8. Define separate pools for the SDIDS, the SDDS, and other data sets.
9. Define separate data and index pools for key sequenced data sets.
10. Create separate SHRPOOLs for production data sets which should be different from other types of data sets, such as, test data sets which do not require LSR pools.
11. Use the QUERY command to monitor the effectiveness of your LSR definitions, and change the number of place holders if you see evidence of waits.
12. Consider using hyperspace buffers.
13. Do not define more than 31 VSAM data sets for one SHRPOOL, as this will cause a U0900 abend to occur.

Tip

When changing your VSAM definition member, keep a backup of your original in case you have to fall back if you have a problem with the definitions to allow your users to continue working in Information/Management until you can correct the code.

3.5 Placement of Data Sets

Place your SDDS, SDIDS, and SDLDS on separate volumes if that is possible. The number of separate volumes is highly dependent on the number of SDDS.

The reason for this recommendation is to reduce concurrent access to the same physical volume during I/O.

Basic I/O performance tuning considerations, such as avoiding heavily utilized DASD or paths, must be followed when positioning your Information/Management data sets. Avoid allocating your databases on the same packs containing the following data sets:

- Page data sets
- Spool data sets
- SMF data sets
- ISPF profile data sets
- Temporary data sets
- Spool data sets
- Sort data sets

3.6 Multiple-Cluster SDIDS

This feature provides relief for the 4 GB limitation on the SDIDS data component. It also allows the user to increase the SDIDS enqueue granularity, as enqueues for each SDIDS cluster are obtained and released independently. As a result, more concurrent processing of the SDIDS can be achieved, improving throughput.

3.6.1 Overview

One objective of this change is to reduce SDIDS contention. To accomplish this, changes to the data generated by assembling a session parameter member and the physical SDIDS record structure are required. This means that migration to Version 1.1 requires re-assembly of all session parameter members (even if not changed) and rebuilding of all SDIDS clusters (even if the SDIDS is not divided into multiple clusters).

Changes to the SDIDS structure include:

- Only 18 (SBCS) and 34 (SBCS or DBCS) key lengths are supported by Version 1.1
- Only advanced compression is supported.

This change enables a database administrator to define a database structure that uses multiple VSAM clusters to contain SDIDS data. Dividing the SDIDS into multiple VSAM clusters reduces the amount of data that is locked when accessing SDIDS data, thus allowing concurrent SDIDS activity to occur, and increases the overall SDIDS storage capacity to more than 400 GB.

SDIDS data can be divided by VSAM record key values. The key of the first record in the first cluster is always assumed to be hex zeros. The key values for the first record in the remaining clusters are specified in a session parameter member via new macros and keywords. The key of the last record in the last cluster is assumed to be less than 32x′ff′.

The database administrator should determine the best way to divide the SDIDS records. This is not a trivial task and there is no automated means of obtaining the ″correct″ divisions. Either the SDIDS analyze utility, or a new GTF user trace point for SDIDS I/O activity monitoring, may be used to collect statistics about

SDIDS activity. Factors to be considered in determining the best way to split SDIDS data include:

- Concurrent update activity on different records in the same cluster
- Update/search contention on records in the same cluster
- The number of records in a cluster
- Amount of read/write activity for each record

Here is an example of a multi-cluster SDIDS, illustrating possible cluster content:

- Cluster 1 - INFO.PROD#01.SDIDS - key range ′0000′x to ′B9FF′x
- Cluster 2 - INFO.PROD#02.SDIDS - key range ′BA00′x to ′BBFF′x
- Cluster 3 - INFO.PROD#03.SDIDS - key range ′BC00′x to ′BFFF′x
- Cluster 4 - INFO.PROD#04.SDIDS - key range ′C00000′x to ′D9D4FF′x
- Cluster 5 - INFO.PROD#05.SDIDS - key range ′D9D500′x to ′D9D5FF′x
- Cluster 6 - INFO.PROD#06.SDIDS - key range ′D9D600′x to ′FFFFFF′x

In this example:

- Cluster 1 contains cognized words that start with lowercase letters
- Cluster 2 contains s-words in the range that includes MASTER_BIT_LIST and LASTENTRYNUMBER
- Cluster 3 contains all remaining s-words
- Cluster 4 contains cognized words in the range ′A.′ through ′RM.′
- Cluster 5 contains all cognized words that start with the characters ′RN′
- Cluster 6 contains all remaining cognized words

3.6.2 Converting from Single to Multiple Clusters

To take advantage of this feature, a database administrator needs to do the following:

1. Analyze the contents of the SDIDS to determine the best way to divide the records. The SDIDS Analyze utility (BLGUT21) output or GTF trace data generated by SDIDS activity can be used to aid in this analysis (see 3.6.2.1, “Dividing the SDIDS” on page 25).
2. Create a session parameter member including the appropriate key range definitions.
3. Define the required VSAM clusters.
4. Load the new SDIDS VSAM clusters.

   A migration utility can be used to copy existing SDIDS data to a new SDIDS structure, or BLGUT1 can be used to build a new SDIDS using an existing SDDS.
5. Update all existing session parameter members that use the new SDIDS to include the new cluster name and key range definitions.
6. Resume normal database activities.

3.6.2.1 Dividing the SDIDS

To divide the SDIDS you need to determine key ranges that help balance the load among the component clusters. One way to divide the SDIDS is to assume that the load is distributed evenly throughout the key range space. The following technique provides a way to split the SDIDS in clusters with the same number of records. Note that this does not guarantee that the I/O is distributed evenly across the component clusters.

1. Find out the total number of SDIDS records with BLGUT21
2. Find out how many records do you want to hold in each cluster by dividing the total number of SDIDS records by the number of clusters desired. Use this number to determine offsets, in terms of number of records, for each range.
3. During a batch processing window (Information/Management should be inactive) run the VSAM utility IDCAMS including a DD statement pointing to your SDIDS (called DDSDIDS in the example) and the following control statements:

   ```
   PRINT INFILE(DDSDIDS) COUNT(1) SKIP(offs-1) HEX
   PRINT INFILE(DDSDIDS) COUNT(1) SKIP(offs-2) HEX
     . . . .
   PRINT INFILE(DDSDIDS) COUNT(1) SKIP(offs-n) HEX
   ```

   Where offs-1, offs-2, offs-n are the offsets, in records, previously determined.
4. From the IDCAMS output obtain the key values corresponding to each range.
5. Code the key values obtained in the session parameter members being used.

3.7 Multiple-Cluster SDDS

This choice should only be considered in a non-Multisystem Database Access environment.

A multiple-cluster SDDS configuration can offer you the following benefits:

- Increased capacity. With a 100-cluster SDDS, the maximum database size is 400 GB.
- Enqueues are independently handled for each cluster.
- The clusters can be positioned on different volumes to cut down on contention.
- Faster backup of databases.
- Faster restore of databases.
- Easier to find space to do a restore for a number of small databases rather than one huge database.
- Faster reorganization of databases.
- Faster access to data due to smaller databases.

You may not need to worry about multiple-cluster SDDS if the total number of records to be stored is about or less than a million records, which can be stored in a single 4 GB SDDS data set. But considering the above mentioned advantages it would be worth implementing.

However, you may consider using a multiple-cluster SDDS even if capacity is not an issue, as there will be a performance benefit if you can place the clusters on separate volumes or separate strings.

If you need a multiple-cluster SDDS, you don′t need to do major changes. However, you must pay attention to backup procedure changes, session-parameter member changes, VSAM resource definition changes, and some syntax changes needed for activating multiple SDDS clusters when running BLGUT1 and BLGUT7 utility programs.

Tip

Remember to use the OTRG=(c,mmm) with BLGUT1 where c is the trigger character and mmm the number of clusters, or the OTRG=(c,mmm) and NTRG=(d,nnn) where c is the trigger character and mmm is the number of clusters making up the source database. d is the trigger character and nnn is the number of clusters making up the target database. If you omit these statements, the utilities will give zero return codes, but the databases will be treated as single cluster databases causing only some of the records to be processed.

It is recommended you use BLX-SP QUERY command together with the IDCAMS LISTCAT (if you can shut down the BLX-SP) to determine capacity in order to carry out tuning. The RESET BLX-SP command is very useful when doing search tuning to reset statistics. *Planning and Installation Guide and Reference*, SC34-4603.

The following section describes some factors you should consider when deciding whether to use multiple SDDS clusters.

3.7.1 Factors to Consider in Multiple-Cluster Configurations

- How large (cylinders/total bytes) is your current SDDS and is it expected to grow significantly?
- Are the advantages of using multiple clusters important?

  The advantages are:

  - DFDSS can be used to do backup and restore to prevent Information/Management downtime, just enqueuing as each cluster is backed up.
  - Some performance improvement. With only one cluster, if a user updates the SDDS, all other users wait until that user completes. Waiting becomes more important when you have many active users (interactive, batch, or API). With multiple clusters, users can access other SDDS clusters while updates are occurring on other clusters.
  - More capacity. One SDDS cluster is limited to 4 GB. Each additional cluster used adds another 4 GB.
  - Reduced time needed for backups. Assuming you have a sufficient number of devices (normally tape drives), you can submit multiple jobs and have each job back up one of the clusters. Since the jobs can run simultaneously, you can reduce the overall time it takes to do a backup.
  - Cut down on the time taken to reorganize the databases by submitting multiple jobs running simultaneously.

The disadvantage is:

– More clusters to manage.

- Take the following steps into consideration before implementing a multi-cluster database:

  – The database is unavailable while running BLGUT7.

  – If you have five clusters and decide to convert to six, you will have to perform all the conversion steps necessary, just as if you went from five to ten clusters, but with only 20% more added capacity. However, if you convert from five to ten clusters, database capacity increases by 100%.

  – Will you have enough DASD to do the conversion? You will need to have twice the amount of DASD tied up while you are converting, all the DASD for the current SDDS cluster(s), plus roughly the same amount for the new SDDS clusters. The old database(s) can be deleted on successful completion of the process. For example:

```
Assume:         Currently a single cluster that uses
                1000 cylinders

Converting to:  4 clusters

Then:           Define 4 new 250-cylinder clusters.
Result:         1000 + 250 + 250 + 250 + 250 = 2000 cylinders needed
                until BLGUT7 completes (the original 1000
                cylinders can be deleted after converting).
```

3.7.2 Determining the Number of Clusters

To help determine how many clusters to use, you can make use of the Access Method Services IDCAMS LISTCAT on your current SDDS cluster(s):

- Look at the HI-USED-RBA value. How close to 4 GB (4 294 967 295) is the HI-USED-RBA?
- How many new records do you expect? Remember that you can search using prefix DATE/ (date entered) to determine how many records per month or year you have created. Also, the BLGUT20 utility provides information on the size of your records.

If your HI-USED-RBA is less than 536 870 912 (half a GB), you will most likely not want to do anything.

If your HI-USED-RBA is less than 2 147 483 648 (2 GB) for the SDDS cluster, depending on how the database is expected to grow, you might benefit from multiple clusters.

If your HI-USED-RBA is greater than 2 147 483 648 (2 GB) per SDDS cluster, you should benefit from multiple clusters in performance and receive 4 GB relief.

If you decide to increase the number of clusters, choose a number that would result in the new clusters being in the range of 536 870 912 (half a GB) to 1 073 741 824 (1 GB):

```
Assume:   Currently a single cluster with HI-USED-RBA
          of 3221225472

          3221225472 divided by 1073741824 = 3

          3221225472 divided by  536870912 = 6
```

You should define at least three clusters. Define four or five clusters if you do not expect much growth in the database. Define six or more if you expect the database to grow over the next several years. Keep in mind that you will have to manage (back up) these data sets.

```
Assume:   Currently a 5-cluster SDDS

          SDDS$01 HI-USED-RBA of 2051235489
          SDDS$02 HI-USED-RBA of 2200929432
          SDDS$03 HI-USED-RBA of 1981778632
          SDDS$04 HI-USED-RBA of 2221225472
          SDDS$05 HI-USED-RBA of 1887331010

          Total number of bytes 10342500035
          (10.3 GB)
```

SDDS$04 is the largest value, so we use it for our calculations:

```
          2221225472 divided by 1073741824 = 2.06

          2221225472 divided by  536870912 = 4.12
```

Doubling the number of clusters to ten would result in each of the ten SDDS clusters having approximately 1 GB of data. Converting to 20 clusters (4 x 5) would mean each cluster would have approximately half a GB. You would need to manage 20 clusters, which might not be worth the effort unless the database is expected to grow beyond 40 GB (4 GB x 10 clusters).

10 to 12 clusters would be a good choice, unless over the long term, the size was expected to grow well beyond 40 GB.

3.7.3 Converting from Single to Multiple Clusters

If you are converting to multiple clusters, follow this procedure:

1. Define the new clusters.

   Reduce the ammount of space for the new SDDS using the same factor you chose for the number of SDDSs. If you are doubling (converting from 5 to 10), use half the number of cylinders for each new SDDS. If you are converting from one to four clusters, each new SDDS will have only one-fourth of the data of the current SDDS, so each new SDDS will need only 25% of the space used by the old cluster.

   Note: Use the same CISIZE, RECORDSIZE, and KEY LENGTH values as your existing SDDS. This will help improve the performance of BLGUT7. You could use the BLGUT20 utility to analyze the SDDS and choose to use a new record size, but BLGUT7 will take approximately twice as long to process.

Tip

Remember the NTRG with BLGUT1 and NTRG with OTRG for BLGUT7 to ensure that all records are processed when using a multicluster database. Also ensure that the clusters are not on heavily accessed system DASD.

2. Update the VSAMDEF.

   - Add a new BLXDSN macro for each of the new clusters.

     Note: All the SDDS clusters can share the same LSR buffer pool. Since you would normally want to have the new clusters use the same LSR pool as the existing SDDS cluster(s), use the existing BLXDSN macro as a model for the new data sets, just add the data set name(s).

   - Reassemble/re-link the VSAMDEF.

     Note: Remember to assemble and link VSAMDEF with a new name then rename the live VSAMDEF to say VSAMDEFB for a backup, then rename the new VSAMDEFN to VSAMDEF. This is to prevent a lengthy downtime if you have incorrectly coded the changes, you merely rename VSAMDEF and your users continue working in Information/Management while you correct the VSAMDEF source in BLXVDEF.

3. Stop and restart the BLX-SP to pick up changes to the VSAMDEF.

   Note: You could use the ADDVDEF command to temporarily add the SDDS clusters. You would not have to stop and restart the BLX-SP; however, you would still need to update the VSAMDEF so that the next time the BLX-SP is stopped and restarted, the correct VSAMDEF would be used.

4. Add TRIGGER to the BLGCLUST macro in your session members BLGSESxx.

5. Update BLGUT1.

   Note: Remember the NTRG parameter.

   - Add OTRG to BLGUT1.
   - Verify that SESS=xx is now the session member you updated or created in the previous step.
   - Verify that the BLGSD DD (if you use it) points to the correct new SDDS data set.

   Note: You do not have to run BLGUT1 after running BLGUT7; however, this is the ideal time to add the NTRG parameter so that BLGUT1 is ready when needed.

6. Using the BLX-SP FREE command:

   - Free the SDIDS.
   - Free the old SDDS.

7. Using the BLX-SP REALLOC command:

   - REALLOC the old SDDS.

     Note: Use the UTIL parameter of the REALLOC command to prevent users from accessing Information/Management prior to running the Index build.

8. Run BLGUT7.

Tip

Be sure to use the NTRIG parameter for the new multicluster database and verify that the NEWSDDS and OLDSDDS ddnames are correct.

9. After BLGUT7 completes:
   - Use the BLX-SP FREE command to free the old single-cluster SDDS.
   - Use the BLX-SP REALLOC command to reallocate the SDDS to allow users to access Information/Management.
   - Use the new session member to access the new multiple-cluster SDDS. Do several searches and display several records to verify the database is behaving normally. You could also bounce the BLX-SP to update the VSAM characteristics then use the LISTCAT command to ensure that there are roughly the same number of records in each SDDS cluster, then comparing the total of these to the total number of records in the original single cluster database.
10. Create or update your backup jobs and procedures.
11. After you complete your tests:
    - Remove the BLXDSN macro for the old cluster in the VSAMDEF.
    - Assemble/link-edit the BLXVDEF source as VSAMDEFN.
    - Rename the original VSAMDEF module to VSAMDEFB. (Delete the old VSAMDEFB.)
    - Rename the new VSAMDEFN to VSAMDEF.
    - Refresh LLA selectively if you are not using a STEPLIB ddname.
    - Bounce the BLX-SP to activate the changes.
    - Delete the old SDDS.

    Note: The change to the VSAMDEF will not be used until LLA has been refreshed (if no STEPLIB is used in the BLX-SP JCL) and the BLX-SP has been stopped and restarted. Leaving the BLXDSN macro in the VSAMDEF will not affect anything.

3.8 Cache/DASD FastWrite

Cache and DASD FastWrite are the combination of hardware and software tools to place data in the storage controller, such as a 3990-G03. Access times to the storage controller (cache), are faster than access to a disk. Access time to the former is four milliseconds as compared to a disk, which can vary from 15 to 30 milliseconds.

The following are the areas where data can be stored and accessed, with their access times:

- CPU memory (LSR buffers) - one millisecond
- Cache (storage controller) - four milliseconds
- Disk - 15 to 30 milliseconds

Memory is the fastest medium, though very expensive. The disk is the cheapest, but very costly in terms of access time. While the LSR buffers are faster than the disk, they are again costly. However, since Information/Management uses storage above the 16 MB line, depending on the customer′s setup for virtual storage use, LSR can be the most effective way for storing Information/Management data.

3.9 Shared DASD When Using MSDA

This type of scenario is not advisable, but you may find yourself in a situation in which you have to face it. In such a case, you must understand the following:

- It is not recommended for optimum performance.
- Multiple and concurrent access to databases from different Information/Management systems can greatly impact performance on all systems, unless all systems minus one access the databases in read-only mode.
- 18 and 34-byte SDIDS key lengths are a good idea in this scenario.
- COGENQ and SRCHLIMIT should be much higher when using shared DASD (in the range of 10 to 20 times more) to complete the record file process to prevent multiple complex-wide ENQ/DEQs per record.

Chapter 4. Memory

Storage requirements for Version 5.1 and Version 6 releases are very similar to Version 4 for interactive users, except for the additional storage required to support the BLX-SP. Eliminating VSAM buffers for each user saves some memory, but unless memory is your most precious resource, you can probably use that saved memory (and more) on relatively large LSR buffer pools.

4.1 Central Address Space Considerations

When the central address space was introduced with Version 5.1, the buffers were moved to a central place instead of being associated with each individual user. Users request the data they need from the BLX-SP service provider, and the latter then handles all the I/Os on behalf of the user. Since most of the data is in the buffers of the server, it is much faster and more efficient to obtain it from the BLX-SP buffers than from DASD.

In Version 4 individual users had a copy of the buffers in their own address space, which resulted in more virtual storage being used in the common service area (CSA). With Version 5.1 and following Version 6 releases most of the I/O is done by the central address space service provider (BLX-SP) on behalf of each user. Each user therefore does not need buffers as large as were required in Version 4.

Depending on the number of buffers and the number of users you may use less or more memory under Version 6.3. The transfer of records from the BLX-SP address space creates an overhead for the memory requirements for the server, other than that, the storage requirements for Versions 5.1, 6.1, 6.2 and 6.3 are similar to Version 4.

4.2 Virtual Storage Constraint Relief (VSCR)

Another inherent difference in the use of storage between Version 5.1 and Version 6 releases, as compared to Version 4, is that most of the buffers in the current versions are stored above the 16 MB line. This offers virtual storage constraint relief, since a lot of storage below the line that was used by Version 4 is now available.

Note: If you see a higher consumption of CSA storage below the 16 MB line, that is probably caused by one of the following:

1. LSR buffers not being used either because they're not defined or the definitions are not correct in the BLXVDEF member. (You must code the LSR buffer pool number for each data set in the BLGCLUST macro.)

   Note: You can use the QUERY command to determine which databases have valid buffer pools defined for them. If you find the database name and only have BLXLxxxx (xxxx will start at 0000 going up sequentially) for the FILE= statement, then you need to add the database to the BLXVDEF source and create a new VSAMDEF module.

2. Too many buffer pools defined.

3. Too many VSAM data sets allocated to Information/Management. For example, too many panel data sets that could be consolidated into a single data set.

4. Allocation of too many NSR buffer pools. Look for BLXVDEF in samplib and check for the BLXNSR macro. The default is 20 NSR buffers. This shouldn't be increased.

4.3 Multisystem Database Access and Memory Usage

There are additional memory considerations if Multisystem Database Access is implemented. This is because it creates the following on all systems that share access to Information/Management data sets.

- A service address space, know as the BIAS
- Subtasks within the BLX-SP
- Subtasks within APPC

In the same way, running APIs from OS/2 or CICS/ESA will involve extra APPC subtasks, and a service address space for the Remote Environment Server.

4.4 Calculating the BLX-SP Region Size

The region size for a BLX-SP server address space can be calculated according to the formula in Figure 6, where the letters in the formula are defined as follows:

P Space required by the local shared resource (LSR) buffer pools. Calculate this as the number of buffers multiplied by the buffer size.

D Space required for data sets that the BLX-SP is required to support.

S Space required for the strings, VSAM placeholders, which are defined for the BLX-SP. Calculate this as the total number of strings allocated multiplied by 1000.

Note: The strings are allocated BELOW the 16 MB line. This can have an impact if you're suffering from CSA memory shortages.

B Space required for the I/O buffers in the BLX-SP. Multiply the maxlrecl for each LSR pool by the corresponding STRNO and add the values for each calculation. This result is added to 4000 to determine the value.

N Space required for non-shared resource (NSR) buffers in the BLX-SP. Calculate this value by multiplying the CI size of each data set using NSR by 3 and adding the results for each calculation.

```
REGION  =   5MB  + P + D + S + B + N


Round this up to the nearest megabyte
```

Figure 6. Calculating the BLX-SP Region Size

The source of the VSAM Buffers member, BLXVDEF, will have the information required to complete the calculation.

Please refer to the *Planning and Installation Guide and Reference*, SC34-4603, for a detailed explanation and an example of the storage calculation.

The above should be used as a guideline only. Since in V6 releases most of the buffers are above the 16 MB line; it does not matter what the user has specified below. This is also dependent on which system USEREXITs, if any, are used to regulate the usage of storage above the line.

MVS/ESA manages the region size allowed based on the REGION parameter coded on the JOB or EXEC statement of the JCL and the EXITs that are being used to manage it if the supplied IBM default exits are being used, and the region is specified as less than 16 MB. That limit applies to storage below the line only. Storage size of up to 32 MB above the line may be used. If a number above 16 MB is specified, then that number is the total limit.

Note: Be cautious about using 0 M for the REGION size, as this tends to only reserve sufficient storage for data sets specified in the VSAM definition member (BLXVDEF in SBLMSAMP data set) which would cause errors when trying to use temporary databases say for archiving or testing.

4.5 Storage Usage by TSPs

There′s no restriction on the number of lines a Terminal Simulator Panels (TSPs) can have in Version 6. Care should be taken before the design and creation of very large TSPs. The TSP is loaded in the user virtual storage before it is executed. The size of the buffer will be the same as the size of the largest TSP processed. The TSP buffer will remain in the user′s address space until the user quits Information/Management. If you′re running very large TSPs, this might increase your memory usage.

Chapter 5. CPU

In terms of CPU utilization, Information/Management Version 6.3 performs very similar to previous V6 releases.

Tip

However, in some cases you may think that Version 6.3 CPU consumption is higher than previous releases. You should not be confused because of it. Some customers have reported higher CPU consumption when:

- Using advanced compression
- Having optimized I/O access by using more data in buffers

In the above cases the reason for higher CPU utilization is that, in fact, more work is performed by time interval, as the system moves from an I/O constrained mode to a CPU constrained mode.

If your customer runs Information/Management in a CPU constrained mode, it may not be wise to recommend the above enhancements, as the outcome will be an increased CPU utilization, with the risk or performance degradation.

If your customer runs Information/Management in an I/O constrained mode, Version 6.3 may help in optimizing the I/O load.

Compared to Version 4, Information/Management Version 5.1 and Version 6 releases accept greater workloads before reaching saturation point.

Clearly, if most of the processor is busy, then CPU can become the key bottleneck. As Information/Management tasks are normally more likely to be heavy I/O rather than heavy CPU consumers, you can assign the Information/Management users to MVS performance groups with higher priority without seriously affecting other system users.

Look carefully at your current performance group thresholds and make sure that your Information/Management users can get sufficient CPU to complete their average transactions. If you have previously tuned your performance groups for Version 4, you should perform a review of your settings after moving to Version 6.

The *Planning and Installation Guide and Reference* tells us the following about Information/Management processor use:

1. In a non-shared environment, the processor cost per transaction with Version 6 (and Version 5.1) is usually similar to Version 4. High CPU tasks such as reports and range searches will use more CPU with Version 6 (and Version 5.1) compared to Version 4.
2. Transactions with high I/O may use more CPU in Version 6 (and Version 5.1) than in Version 4.
3. Processor costs for shared DASD in Version 6 are higher due to the cost of GRS and buffer invalidation.

Planning and Installation Guide and Reference, SC34-4603.

5.1 CPU Considerations for Version 6 Releases

Information/Management Version 6 releases and Version 5.1 use a different approach for VSAM I/O than prior versions did. Instead of each user directly dealing with VSAM, the user requests for VSAM access are coordinated through the BLX-SP. This provides significant performance advantages (two to three times maximum throughputs) due to improved VSAM buffering through the BLX-SP LSR buffers. However, there is a cost. That cost is the extra CPU time required for each VSAM input or output. For many transactions that only do a few VSAM I/Os, the cost will be negligible. For transactions that do many (10-200) VSAM I/Os, the increase in CPU usage varies from 10 to 100% depending on the amount of I/O. For a few transactions types (notably large range searches or sorts requiring range reads of the SDIDS) the extra CPU usage can be several times that of Version 4.

At throughputs near Version 4 saturation, Version 6 and Version 5.1 used about 20% more CPU time for our representative workload. The CPU cost per transaction at higher workloads in Version 6 and Version 5.1 continues to drop and approaches the best that Version 4 was able to obtain for the same workload.

However, if you are moving to Version 6 from Version 4, expect some CPU usage increase. You should attempt to eliminate any very large range reads of the SDIDS due to inefficient search criteria or sorting of large search results by inappropriate fields. The most likely source of such range reads are RNID/ searches or sorts of search results lists by RNID.

If you′re at Version 4 and in a CPU constrained environment, carefully consider the CPU usage implications before moving to Version 6. By reviewing your most frequent search criteria and sort fields you can avoid range reads of the SDIDS and prevent or reduce potential CPU usage increases after installing the new release.

5.2 Comparing CPU Usage in Different Versions

This section compares Information/Management Version 4 and Version 5.1 CPU utilization based on measurements and analysis done after delivery of Version 5.1. Performance of Version 6 will be similar to Version 5.1 in a non-shared DASD environment. It is included here for your reference.

The chart in Figure 7 on page 40 plots total system (SYS) and total task control block (TCB) CPU seconds per significant Information/Management transaction from a TPNS driven workload used against both Version 4 and Version 5.1. The SYS and TCB numbers came from RMF and were divided by the throughput in significant Information/Management transactions per minute and plotted against that throughput.

You will need to understand the following items to interpret the chart:

- Representative workload here means 30 users all doing a similar loop of interactive actions in Information/Management. Filing (new), displaying, and filing (updated) records, plus searching and scrolling through search results lists were the main types of significant Information/Management transactions (significant means affects the SDDS or SDIDS).

- The representative workload we defined contains no BATCH, API, sorting, or other complex automation (TSPs, user exits, etc.) that some customers use. CPU costs for some transactions not included in the workload can be substantially higher in Version 5.1 than in Version 4.
- In addition to the significant Information/Management transactions, the workload also included many trivial Information/Management transactions (which did not impact the SDDS or SDIDS) which typically involved simple panel changes or returns to the primary option panel. On the average there were about four trivial transactions per significant transaction. Even though the graph is plotted in terms of significant transaction, the CPU costs shown include all trivial transaction costs.
- Measurements were made starting with relatively low throughputs. Then, the TPNS user think time was decreased for the 30 simulated users in steps until Information/Management simply could not handle more transactions per minute.
- Version 5.1 LSR buffers included all of the index components of the SDDS and SDIDS, but very small numbers of LSR buffers for the data components of those critical data sets. (Version 5.1 was able to handle throughputs approaching 300 significant transactions per minute on other runs which had larger buffers with no significant change in the CPU cost per transaction.)
- Both Version 4 and Version 5.1 measurements were made on the same second-level system. The second level system was running under VM on a 3090 400E. The second level target (with Information/Management on it) in both cases was only allowed access to two of the four CPUs, so 100% means all of two processors.
- TPNS was run from a different second level under the same VM first level and was restricted to one processor. Runs were made off-shift, but no controls existed to prevent other users on other second levels under the same VM from using resources. Although it′s true that this can cause fluctuation in the answers, sufficient runs were made to make us believe that the impact of other users is insignificant.

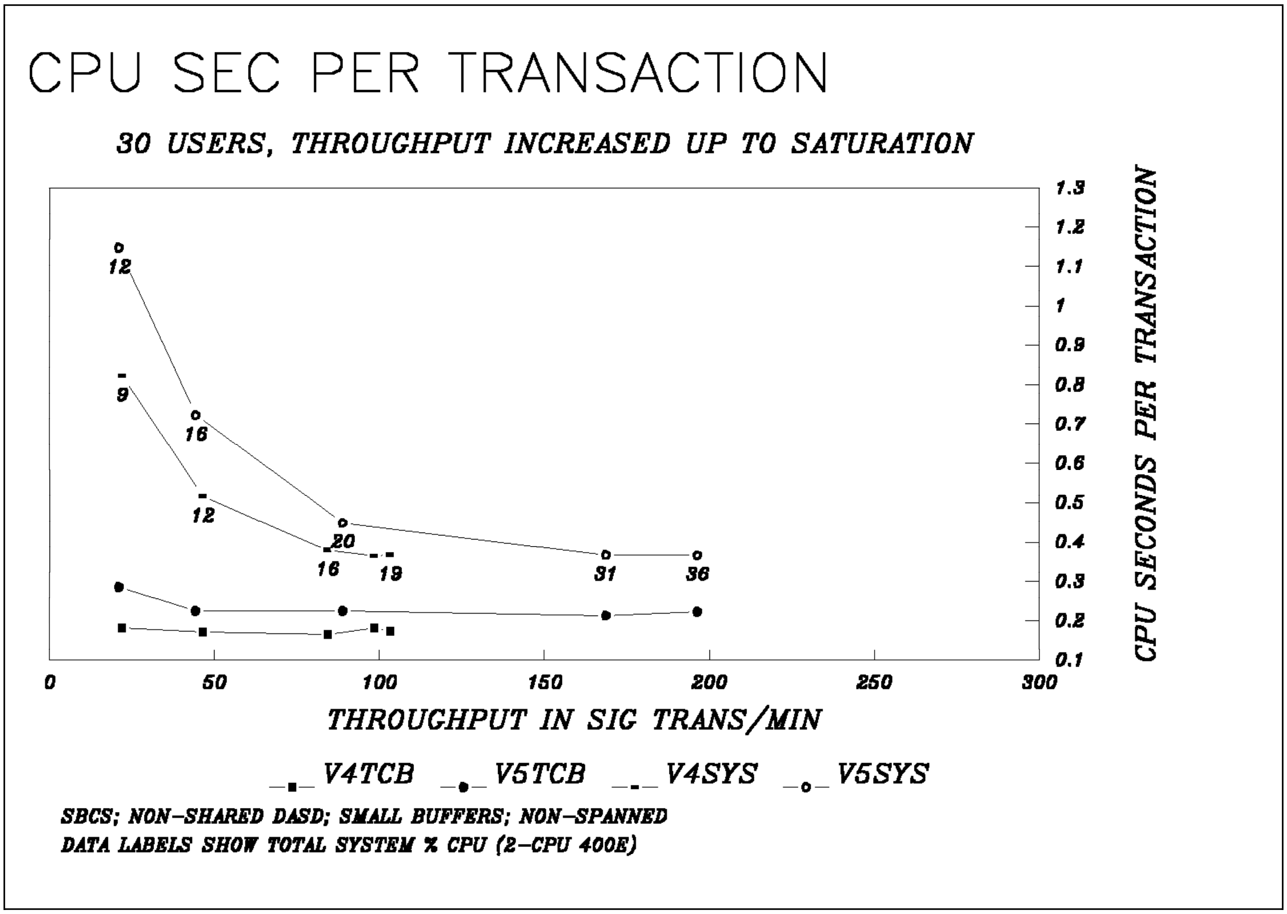

Figure 7. Comparison of CPU Use for Version 4 and Version 6

Significant points from the chart shown in Figure 7 include:

- Version 5.1 CPU costs per significant transaction are higher than Version 4 in this representative workload. The percentage increase is greater in low throughput situations and decreases at higher throughputs.
- Version 4 was unable to process much over 100 significant transactions per minute, but Version 5.1 was able to handle about 200 per minute, becoming more CPU efficient at throughputs above those possible with Version 4.
- Comparing CPU seconds per transaction at throughput range of about 80 to 90 significant Information/Management transactions per minute range shows that TCB times are increased by about 35%, but SYS times are only increased about 18%. Users should be careful not to compare TCB times and think things are worse than the total SYS times show.
- At lower throughput rates the SYS CPU per transaction increases for both Version 4 and Version 5.1, but at such lower workloads, the net effect of Information/Management on the total system CPU use becomes less important.
- At throughput rates above those possible in Version 4, Version 5.1 CPU cost per transaction continues to fall to about 0.4 seconds per significant transaction, which was the best that Version 4 could do.
- The numbers on the two SYS lines show overall percent utilization of the target (two-processor) system. Thus at throughputs near its throughput saturation, Version 4 did not use more than 19% of the two processors.

While Version 5.1 at similar workloads was using as much as 18% more than Version 4, its net effect on the overall CPU utilization was only about 1-2% of the total processor power of the target.

5.3 Doing More Work?

Version 6 and Version 5.1 can allow you to handle more work, much faster than Version 4, but you will need CPU power available to handle the extra throughput.

Information/Management Version 4 is usually limited in how much CPU power it can use. Its maximum throughput is normally limited by the enqueue contention on the SDDS and SDIDS. Version 5.1 provides significant enqueue relief (up to three times the maximum throughput compared to Version 4), so you may easily increase the amount of work that you can do with Information/Management. Even for transactions where the CPU cost is the same in Version 4 and Version 5.1, if you are doing 30% more transactions per minute, the CPU cost will increase by 30%.

5.4 More about CPU and Type of Transaction

All known increases in CPU use are related to VSAM I/O needs of various transactions. Several general categories of transactions can be considered based on VSAM I/O counts:

- VERY LITTLE: Many transactions do not need VSAM I/O. These should not show any increase in CPU needs and may in some cases use less CPU.
- SMALL: Most displays or simple searches will only do a few (less than 10) VSAM I/Os and will tend to have CPU costs similar to Version 4.
- LARGER: Transactions that may do more than ten but less than a few hundred VSAM I/Os may typically experience CPU cost increases of from 5 to 40%. Transactions which may fall into this category include more complicated searches, many sorted searches, small reports, scrolls, and filing of new, changed, or deleted records.
- Apply Information/Management fix UW06240 to solve APAR OW04230 (improves performance from the Information/Management side).
- Consider using TCP/IP connections and Windows for your low-end workstations.
- MASSIVE: Some specific situations where certain transactions may require truly massive amounts of VSAM I/O (tens of thousands to millions of VSAM I/Os) such as with very large range searches or sorts of large SRLs on fields with huge numbers of different values. These types of transactions have always been bad; Version 5.1 does not read any more of these records than before, but since they are almost entirely focused on that massive VSAM I/O, their CPU needs can be several times that required in Version 4. For more information about massive range reads and how to avoid them, please refer to Appendix A, “Searching and Sorting Flash” on page 81 for an extensive discussion of search parameters.

5.5 Summary of CPU Considerations

When using Version 6 and Version 5.1, an increase in the utilization of the CPU will be noticed compared to Version 4. The increases are due to the more CPU-intensive I/O processing that uses the central address space to speed up the I/O via LSR buffering. Thus, the additional cost will vary by transaction type, or, more specifically, by the amount of VSAM I/O caused by a given transaction.

For representative workloads such as the one analyzed here, the increase in CPU use at the total system level will probably be a relatively small percentage of the total CPU power available. Information/Management is rarely able to use a large percentage of a processor complex before it limits its own throughput due to SDIDS enqueue contention.

There are individual transactions that can do massive numbers of I/Os, which will use substantially more CPU in Version 6 and Version 5.1 than in Version 4. In these cases, tuning and avoidance techniques can sometimes help.

Please refer to Appendix A, “Searching and Sorting Flash” on page 81 for an extensive discussion of search parameters.

Chapter 6. Impact of new Information/Management Features

This chapter focuses on the new features that have been introduced since Information/Management Version 6.1. Most of these can be utilized to enhance performance. However, some of them can and do impact performance negatively in certain circumstances. Hence, the circumstances that can cause degradation are discussed here, with advice given on how to avoid them.

6.1 Spanned Records No Longer an Issue

The need for VSAM spanning of the SDIDS is eliminated when using a new SDIDS key length of 18 or 34 bytes. The new key lengths also allow any number of records to be stored in a database without the need to change the SDIDS record size. SDIDS records that were spanned can now take advantage of the local shared resources (LSRS).

6.2 SDIDS 18 or 34-Byte Key Lengths

You should consider using the 18 or 34-byte SDIDS keys provided in Version 6.3 of Information/Management to eliminate the need to span VSAM records and to significantly improve throughput and performance.

Figures taken from actual customer installations show a performance gain of 30%.

6.3 Multiple SDIDS Data Sets Support

This feature provides relief for the 4 GB limitation on the SDIDS data component. It also allows the user to increase the SDIDS enqueue granularity, as enqueues for each SDIDS cluster are obtained and released independently. As a result, more concurrent processing of the SDIDS can be achieved, improving throughput.

6.4 SDDS Size Extended

The maximum storage capacity of SDDS is extended from 20 to 400 GB (5 to 100 clusters). Previous support only was for 1 or 5 SDDS clusters at 4 GB each.

Note: The current limit of logical SDDS records is 16 777 215 for 16 or 32-byte SDIDS key lengths.

6.5 Advanced SDIDS Compression

Advanced SDIDS Compression is optional in Version 6.3 and standard in Version 1.1.

Advanced SDIDS Compression can improve performance, even when dealing with non-spanned records.

When the SDIDS is spanned, Advanced SDIDS Compression can improve performance in high SDIDS contention situations, as well as allow more records to be stored in the SDDS before the SDIDS again becomes spanned.

These benefits have been shown to be substantial (increased maximum throughputs of 20% for non-spanned SDIDS and 100% or more for spanned SDIDS) in development lab measurements. Advanced SDIDS Compression should lead to lower SDIDS enqueue contention, and improved response times in heavy workloads with no serious CPU overhead.

New SDIDS records with key lengths of 18 and 34 bytes can also benefit from using advanced compression, since compression reduces the size of SDIDS records.

Note: When using a 16 or 32-byte key length, you can define the SDIDS as SPANNED before the number of records in the SDDS actually goes over the 261 000 limit, thus going spanned. This would prevent message BLG19509E informing you that the maximum number of records allowed in the database has been exceeded. This could mean long downtime if it happened during the day as you would have to delete and define your SDIDS and rebuild it using BLGUT1, which if you have a database of about 2000 cylinders could take up to 2 hours.

If you want to use the Advanced SDIDS Compression, then you can do so in any of the following ways:

- If you can run BLGUT1, then just specify COMPRESS=ADVANCED in the parameter list and the entire new SDIDS will be compressed using this new methodology.
- If you plan to start with a new, empty SDDS, then you can run BLGUT2 against your newly defined SDIDS before adding any records to your database. If you select the COMPRESS=ADVANCED option under BLGUT2 for an empty SDIDS, then when you start to use that SDIDS, all records will be compressed when they are first created during normal operation.
- If you plan to start with an existing SDDS and do not plan to run BLGUT1 to create a new SDIDS, you can still use BLGUT2 to specify COMPRESS=ADVANCED for the existing SDIDS. Your SDIDS records will gradually become ADVANCED compressed as they are used the next time that the BLX-SP is started. It may take hours or even days for most records to become compressed using this gradual method, but most of the benefits will be realized very quickly after the first records are filed.

The diagrams in Figure 8 on page 45 illustrate that to gain the full benefit it is better to run BLGUT1. However, BLGUT1 takes longer to run than BLGUT2. Conversely, BLGUT2 can be run quickly to enable the feature, but will require a longer time before the benefits will be realized.

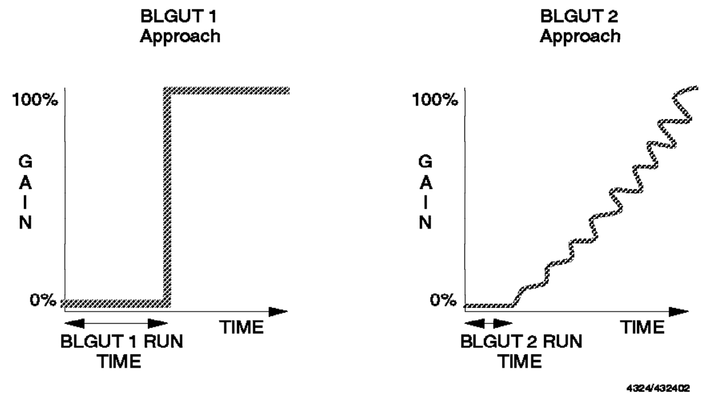

Figure 8. Implementation of Advanced SDIDS Compression

If you do not wish to use Advanced SDIDS Compression, then you can simply use your Version 4 or Version 5.1 SDIDS directly in Version 6. You do not need to run BLGUT1 or BLGUT2. The default compression technique is BASIC, which is exactly like that used in Version 4 and Version 5.1.

It is possible to make the master bit list in the SDIDS fit into a single VSAM control interval (CI) as this record, the longest in the SDIDS data set, is dramatically shortened by the enabling of Advanced SDIDS Compression. This alone can save processing time as the master bit list will fit into buffers if it does not span more than one CI.

There are four ways that a performance gain is achieved from enabling Advanced SDIDS Compression:

1. Non-spanned SDIDS databases show increased throughput which is estimated to be between 5% and 20%.
2. Spanned SDIDS databases show increased throughput, with the greatest gain possible in situations where the master bit list spans multiple CIs.
3. Individual response times can be improved. The largest gain here is where a search would read several long records from the SDIDS in the absence of Advanced SDIDS Compression. Updates can also show an improvement, whereas creates and deletes are likely so show little difference.
4. With a large database at peak activity, response times will be shorter. Some of this gain will be an indirect benefit from reduction in delays due to greater buffer utilization.

Once any SDIDS has been used with Advanced SDIDS Compression you would have to delete and define then rebuild the SDIDS using BLGUT1 specifying COMPRESS=BASIC in order to use it again with Version 4 or Version 5.1. For this reason, it is recommended not to implement this until your Information/Management Version 6.1 database is stable (same applies to Version 6.2 and Version 6.3).

Tip

If you archive (or purge) 10% or more of your database records this can seriously affect the compression of your SDIDS. This is because the 1′s compression algorithm will be less effective as there will be shorter blocks of 1s in your master bit list. Under certain circumstances, this might even cause your SDIDS records to become spanned, severely degrading performance. To avoid this problem, allow for the ″gaps″ in the master bit list to be filled with new record creations by allowing some time between archives (or purges) and maintain the number of records archived (or purged) at about 5% of your total record number. Ideally, you should balance the number of record deletions and creations to maintain your database size constant, unless you rebuild your SDIDS to do a cleanup.

6.6 Queued Notification Mail Processing

Notification mail can be queued to the BLX-SP as an alternative to processing mail immediately in the user′s address space. Use of a BLX-SP queue can improve the performance of immediate notification since not all the processing is done at the time the record is filed.

Some of the processing can be done by one or more dedicated mail address spaces (another interactive or batch Information/Management session) that get the mail from the BLX-SP and process it.

6.7 Multisystem Database Access

Since Version 6.3 Information/Management includes a new performance function related to Multisystem Database Access. When a BLX Service Providers (BLX-SP) detects a failure in an APPC conversation with a partner, it will make one attempt to reallocate that conversation. If that attempt fails, the BLX-SP will wait at least 5 minutes before making another allocation attempt. This process is repeated until allocation succeeds. Performance is enhanced by eliminating the many futile allocation attempts made by previous versions of Information/Management in this situation.

Also, a timer function is implemented that detects when Multisystem Database Access is hung waiting on a response from APPC/MVS. The wait terminates after 30 seconds, the conversation deallocates and an error message is issued, allowing the process to continue. The conversation is then reallocated just as it would be in the event of any other communications failure. Performance is enhanced since users won′t have to queue up needlessly behind a hung user waiting on an enqueue.

No loss of data integrity occurs in either of these scenarios.

Information/Management Version 6 releases provide the capability to allow multiple BLX-SPs on the same or different MVS/ESA systems to share common Information/Management data sets through the use of Advanced Program-to-Program Communication (APPC) and Global Resource Serialization (GRS).

Information/Management V6.3 Installation as an Enterprise-Wide Server Web Connector, E-Mail and Remote API Support explains how to implement this feature, reintroduced with Information/Management Version 6.1. However, care needs to be taken with this and the consequences from a performance point of view need to be clearly understood. Performance benchmarks have shown that Information/Management performs more efficiently in a non-shared situation. This section is intended to explain the factors that will affect performance in a Multisystem Database Access environment and the areas in Information/Management, APPC, and GRS that can be tuned. The diagram in Figure 9 illustrates a sample configuration.

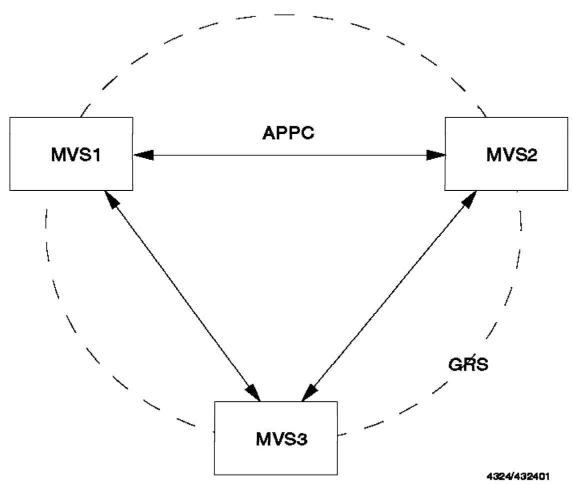

Figure 9. How Multisystem Database Access Is Implemented

The factors that can influence performance in a shared situation are as follows:

1. The number of sharing systems in the GRS loop. This will have a direct bearing on the time it takes to perform the polling around the loop which will in turn affect end user response time.
2. The GRS parameters RESMIL and ACCELSYS will directly affect the performance of the shared Information/Management systems.
3. The change activity on the VSAM data sets. With a larger volume of change activity, there will be a greater number of buffer invalidations required to maintain data integrity. The size of the change is less significant than the number of changes. Use the COGENQ parameter to help reduce the number of fields that are cognized to the SDIDS per enqueue.
4. The first filing transaction after a restart of BLX-SP in a shared situation will always be significantly longer than the second and subsequent. This is because the Remote Environment Server (RES) associated with each BLX-SP is started and initialized by APPC at this time. If a TSP were started immediately after all connected BLX-SPs were up and running, then this would alleviate this problem due to the panels having been loaded into LSR buffers.
5. The use of multiple-cluster SDDSs is not recommended in this environment.

6. There is a CPU overhead in handling the buffer invalidation process that is necessary for data integrity.
7. Large searches and large reports, and their corresponding high read activity can be important in this environment. This is because of additional Version 6 enqueues, which in a shared environment must be global.

If you plan to use Multisystem Database Access, you should consider the following items:

- There are significant CPU costs for Multisystem Database Access. The use of subtasks under the BLX-SP to communicate VSAM buffer invalidation information to all other sharing BLX-SPs through APPC adds to the CPU costs for a given workload. In developmental measurements using ″representative″ workloads, this increase was in the vicinity of an extra 30% for overall CPU utilization at the system level.
- Versions 5 and 6 both use extensive interleaving of enqueues; that is, enqueues are released and re-obtained frequently during transactions. In Version 5.1, that is not very expensive because data set sharing among BLX-SPs is not allowed so all enqueues are local. In Version 6, Multisystem Database Access uses GRS (or an equivalent) to communicate enqueues among sharing systems. To reduce the number of enqueues that must be globally obtained, Version 6 includes some parameters which allow you to restrict the interleaving of enqueues. You should consider setting these for best performance if you are using Multisystem Database Access. For details, see the Performance chapter of the *Planning and Installation Guide and Reference*.
- GRS (or an equivalent) also contains parameters (see RESMIL and ACCELSYS for GRS) which must be balanced to provide the best performance for your multiple-system situation. Lowering RESMIL can cut enqueue delays, but will increase CPU utilization.

With sufficient CPU resources and reasonable tuning, Version 6 can perform much better than Version 4 in shared situations. Using a ″representative″ workload in a complex of three sharing systems, Version 6 was able to double maximum throughput compared to Version 4. Average response times were also much better. Individual response times will vary and will not always be better in Version 6. Sharing data sets in Version 6 is a significant performance impact; the same workload that doubled throughput in Version 6 (compared to Version 4) would be approximately doubled again in Version 6, without sharing.

Customers who operate in a CPU-constrained environment are not encouraged to use Multisystem Database Access without carefully considering the CPU implications.

Version 5.1 customers who have taken advantage of the increased maximum throughputs of Version 5.1 should also consider carefully the performance impacts of Multisystem Database Access, since Version 6 (when sharing data sets among BLX-SPs) will not perform as well as Version 5.1, which cannot share data sets.

6.7.1 GRS/MiMs

Each resource is identified by a major name (QNAME) and a minor name (RNAME). Information/Management requests the correct level (local or global) of resource serialization that it needs, but its requests are sometimes overridden by your own GRS or MiM parameters.

Version 5.1 and 6.1 customers are the most likely to be affected by incorrect settings, but Version 4 customers may also benefit by correcting their GRS settings. Because the correct settings changed for each version, your own overrides remaining from past versions may adversely affect Information/Management performance when you migrate to a newer version.

6.7.1.1 Using IBM's GRS

GRS automatically honors the local versus global requests by an application, unless you override the application's request in SYS1.PARMLIB in the GRSRNL member. That member allows you to specify overrides by QNAME and RNAME. You need to be sure that any overrides are correct.

No matter what version of Information/Management you use, verify that there are no QNAMEs beginning with BLX or BLG in your GRSRNL member. If you are a Version 4 customer and you are not sharing data sets, you might slightly improve performance by using the GRSRNL member to override the QNAMEs beginning with BLX and BLG, making them local. You must be sure, however, that those data sets will *never* be accessed from other systems; otherwise, the data sets may be corrupted.

6.7.1.2 Using Legent's MiM

You must tell MiM how to handle resource requests. MiM has three members that you should verify:

- The MIMINIT member uses the PROCESS= option so you can specify how you want resources to be treated.
- The MIMQNAME member lists QNAMEs and specifies several parameters for each. Parameters GDIF and SCOPE define how to generally handle each QNAME, but parameter EXEMPT enables you to override that general specification for individual RNAMEs as defined in an GDIEXMPT member.
- The GDIEXMPT member enables you to override the default or more general specifications with specific local versus global definitions by RNAME.

If you are an Information/Management Version 4 customer, be sure the MiM parameters force the QNAMEs beginning with BLX and BLG to be global. If you can guarantee that you are not sharing data sets across systems, then you may see small performance improvements if you use the defaults or set the Information/Management resources to local.

If you are a Version 5.1 customer, you should always force MiM to make all QNAMEs beginning with BLX and BLG local. (The one exception is BLXCAS, which should be global.) You should be sure that any prior global overrides are not remaining.

If you are a Version 6.1 customer, and you are not sharing among multiple systems, then the same MiM setup recommended for Version 5.1 customers applies to you. If, however, in Version 6.1, you are sharing among multiple systems, then you must be sure that MiM treats all QNAMEs beginning with BLX and BLG as global.

If you are running two or more Information/Management installations in the same set of linked systems and they have different serialization needs, then you should use overrides for GRS or MiM to keep any RNAMEs that should be local from being handled globally. Where the RNAME includes a parenthetical expression such as data_set_name(record_id), you can use a wildcard to cover all record_ids.

In summary, serious performance problems can occur if requests that should be local are overridden as global. Data set integrity problems can also occur if requests that should be global are overridden as local. The global versus local settings for Information/Management resources vary by version and by your choice of serialization tools. If you are having performance problems or if you are migrating to a new version of Information/Management, it's important to be sure that your GRS or MiM specifications are correct.

Tips

- When using Global Resource Serialization (GRS) on multiple systems, ensure the GRS RESMIL parameter is set to 1 on all systems.
- It is not recommended to run Multisystem Database Access on a single MVS machine, as performance will be impaired. The reason for this is the process of buffer invalidation introduced above. This method of maintaining data integrity requires that the BLX-SPs communicate through APPC. When a user issues a file request while connected to one server, then before his or her request can complete, the other server or servers must invalidate their buffer containing a redundant copy of the record.
- Multisystem Database Access is designed to run on two or more separate MVS/ESA images using multiple servers, with two or more engines sharing the workload.

6.8 BLGUT3 and BLGUT4 Utilities

These database utilities now offer improved performance. Multiple updates in the offloaded log data set (reflecting changes since the most recent database backup and SDLDS offload) are pruned by the BLGUT4 utility so that a record being restored will be processed only once or twice no matter how many times it was written to the SDLDS.

For example, if a record is created, it is processed once. If a record is created and updated once, it is processed twice. If a record is updated multiple times, only the last two updates are processed.

6.9 Automatic Log Save

Installing Information/Management Version 6 Release 1 for Client/Server, GG24-4323, provides details of how to implement this new function. There would be two primary reasons for implementing Automatic Log Save. These are:

- To establish a read-only database that is to be used primarily for reports. This would be intended to protect the performance of the prime database.

- To create a Hot Standby system, again read only in a normal situation, that could be used as an alternate system if required. This would need to be on a separate DASD string, or even on a remote site, in which case it would also serve as the Disaster Recovery System too.

The consequences of running Automatic Log Save are as follows:

1. Source impact at capture time.

   At time intervals specified in the LOGSAVE record on the source database, a timer will trigger the capture process. There is an enqueue held on the source SDLDS that will impact filing activity on the source system for the duration that it is held. This enqueue will be influenced by tuning the RECS= parameter in the TSP BLGDUMP1 to reduce the duration that the offload process will enqueue.

2. The target database needs to keep up with the source database.

   The mirroring activity is no different to an interactive user in I/O and enqueue delays. Clearly, if the target database is on a remote machine, then network delays will cause it to lag behind to some degree. The delay may only be a problem in a situation where the target database is seen as a critical Hot Standby. In this case, the LOGSAVE record would be set to have more frequent mirroring activity, and the NetView File Transfer program could be used to speed up the mirroring, as described in *Information/Management V6.3 Installation as an Enterprise-Wide Server Web Connector, E-Mail and Remote API Support*.

3. The target database competes for CPU with the source database if they are on the same machine.

   If the two databases are on the same system, the source system should be given a higher priority to avoid the two competing on equal terms for processor time.

Tips

- Automatic Log Save can be used to move the overhead of running RFTs to a separate database.
- Automatic Log Save can be used to assist in establishing a Disaster Recovery database.
- The history values are not the same in the production database and in the backup database, because duplicate records are eliminated in the LOG.

6.10 DB2 Extract Facility

This is essentially a variant on Automatic Log Save, in that you are creating a downstream copy of your database. Again, the target database will be operating in a fashion described above for the Automatic Log Save target database running on the same machine, so the performance considerations are the same.

Tip

DB2 Extract Facility can be used to move the overhead of running reports away from your production database. Reporting would also be easier if you have SQL programmers.

6.11 Multiple Central Address Spaces

You can have more than one BLX-SP with since Information/Management Version 6. The delimiting factor in terms of the number you can have is essentially processor capacity. As in terms of what the Information/Management product supports, you could define BLX0-BLX9, and BLXA-BLXZ or any other four character names in your subsystem name table. In practice, it is unlikely that there will be more than two or three. The implications of sharing data sets within the same MVS/ESA environment have been discussed in 6.7, "Multisystem Database Access" on page 46, so in all likelihood, two BLX-SPs are likely to have independent data sets, possibly one address space acting as a server for their production database and one for the test database. From a performance perspective, this setup would entail minimum overhead. If you do not have to bounce the BLX-SP, then one BLX-SP is the way to go.

6.12 API Performance Considerations

There are several clients that are supported by Information/Management:

HLAPI/NT This is the Windows NT-based client. It uses TCP/IP to communicate API requests from the workstation to the host database.

HLAPI/UNIX This is a multiplatform client. Information/Management supports AIX, HP and Sun Solaris clients. All can use TCP/IP to communicate API requests from the workstation to the host database. The client for AIX can also use APPC.

HLAPI/2 This is the OS/2-based client. It uses APPC to communicate API requests from the workstation to the host database.

Details of how to install and exploit the above APIs can be found in *Information/Management V6.3. Network Computing and Workstation Interfaces*, SG24-4574.

HLAPI/CICS This is the CICS/ESA-based client. Again APPC is used to channel API requests across an LU 6.2 link. Details of how to implement this can be found in *Information/Management V6.3 Installation as an Enterprise-Wide Server Web Connector, E-Mail and Remote API Support*, SG24-4573.

The most important things that will affect the performance of your CICS/ESA region will be the following CICS/ESA parameters:

- MXT - This system initialization parameter limits the number of concurrent tasks in the CICS system.
- AMXT - This system initialization parameter limits the number of tasks being handled concurrently by the CICS dispatcher.

- CMXT - This system initialization parameter limits the number of tasks handled by the CICS dispatcher within classes defined with the TCLASS operand.

  If you choose to define all BLM transactions to a TCLASS, and you set a limit of 10 for this class, then if an attempt is made to start an eleventh BLM transaction, that user will hang until one of the ten slots becomes free.

- Number of sessions - When you define your LU62 session from CICS/ESA to APPC/MVS, then you specify a value for the number of sessions that you wish to allow on the link. This parameter limits the number of concurrent HLAPI/CICS users. If you set this to 5, and a sixth user tries to start an HL01 transaction from CICS, then that user will hang until one of the existing HLAPI/CICS users issues an HL02 transaction to free up a session.

  Please note that if you set the number of sessions to 5, and you have five HLAPI/CICS users present, then this will use up five of the slots available for active tasks, as there will be a BLMK task associated with each HLAPI/CICS user.

  It is important to understand that if you plan to run HLAPI CICS in the same CICS region as an existing CICS application, then you should review the values for MXT and AMXT in your system today. If your CICS region is busy before you install HLAPI/CICS, and you allow too many HLAPI/CICS tasks to run, then this could impact the existing application.

Therefore, if HLAPI/CICS is being run on an CICS/ESA region that is running other applications, then thought needs to be given regarding performance. Please see *CICS/ESA Performance Guide*, SC33-0659, for further information.

Tips

- If you run HLAPI/CICS in a CICS/ESA region that is running other applications, then consider putting the HLAPI CICS transactions in a transaction class of their own, with a CMXT limit.
- The number of sessions on the LU 6.2 link will limit the number of concurrent HLAPI/CICS users.

6.13 ISPF Client/Server

The ISPF client/server feature is available if ISPF 4.1 is installed. This feature is also known as the graphical user interface (GUI) or the ISPF GUI.

This is a very promising feature of ISPF and worth looking into. Combined with appropriate PC hardware it provides a very modern and appealing interface.

Note that it will not perform as fast as a regular TSO user on a 3270 terminal or emulator. This is due to the additional costs involved in running the ISPF GUI:

- The ISPF style panels have to be mapped to PM equivalents on the workstation each time the panel is used.
- There is the overhead of the work done by APPC.
- There is the overhead of the code that runs on the PS/2.

Tips

- If you choose to run the GUI, then it is best to have a 486-based workstation with a fast hard disk.
- Apply ISPF fixes for APARs OW03965 and OW03966 to improve performance.
- Apply Information/Management fix UW06240 to solve APAR OW04230 (improves performance from the Information/Management side).
- Consider using TCP/IP connections and Windows for your low-end workstations.

6.14 List Processor

This is an area where significant improvements have been made to reduce processor costs. This has been achieved in the following areas:

1. Starting the list processor.

   The savings here vary according to the size of the list. A potential 60% savings is possible with very large (for example, 600 rows) single-column lists. Savings can also be made with very large (for example, eight columns of 3200 rows) multiple-column lists, and medium size (for example, three columns of 300 rows).

2. Filing a record with large amounts of cognized data.

   The largest gains will be in situations where a small proportion of the data has been modified, but checks are needed to see what fields are cognized. The greater the amount of cognized data, the greater the CPU saving compared to Version 5.1.

3. Saving a list.

   A large list with a small proportion of changes will give the greatest improvements.

6.15 APIs

There have been several areas that have improved the performance of the API. These are:

- Dynamic PIDTs

  These can sometimes lead to smaller PIDTs to be processed, and can result in fewer PIDT reads for certain types of activity. The elimination of some PIDT accesses due to dynamic PIDT can yield significant performance advantages for some workloads.

- Cached API Tables

  When an API program repetitively uses a table, it is held in cache instead of reloading it multiple times. Before this was made available, some API programs spent a large proportion of their run time waiting for physical reads.

- Searching

  In earlier versions of Information/Management, if an API program issued a search request, then the search would behave differently to an interactive user. Where an interactive user would read in only sufficient records to fill the 3270 screen, an API program had to retrieve all SDDS records matching the search criteria. This clearly impacted the performance of API programs. Now the API search behaves the same way as an interactive user, thereby improving performance.

6.16 WORDFIX and S-Word Cognize

The SDIDS will grow in relation to an increase in fields that are either p-word cognized or s-word cognized. With this increase, there will come a performance cost. Conversely, reducing the number of fields that are cognized will reduce the size of the SDIDS and improve performance,

This benefit can be achieved by taking one of the two following approaches:

WORDFIX Use WORDFIX to uncognize fields that were previously cognized. BLGUT1 needs to be run to reap the benefits.

Optional S-Word Cognize It is now possible to specify s-words that users do not want cognized, while leaving the p-word cognizable. The new view internals can be used to identify which s-words should have cognize removed (see the Customer Study section for details).

The effect of both these changes is a reduction in the size of SDIDS, which will improve the buffer efficiency so that the reduced I/O will reduce response times.

6.17 REXX

REXX programs can run in an interpreter mode or in a compiled mode, but performance-wise it is better to run compiled REXX programs.

The performance improvements that you can expect when you run compiled REXX programs depend on the program. Better performance improvements are achieved when the program performs large numbers of arithmetic operations with default precision. Performance gains may go from two to six times faster in many cases, with the possibility of jumping over ten times faster when performing arithmetic operations.

Chapter 7. Customer Performance Study

Sometimes you run into a situation in which a customer tells you that they are not fully satisfied with the overall performance of Information/Management. Even though we have been discussing performance issues, factors, areas and parameters, we haven't described an overall technique to analyze, address and identify potential performance problems, before you start performance tuning tasks.

This chapter can help you understand how to approach performance analysis and tuning by means of a reference example taken from a real customer. The technique used leads you to the right information by paying close attention to the following tasks:

- Collecting installation-dependent definitions
- Analyzing the application
- Analizing SDIDS data sets for p-words
- Analyzing report data sets for range searches
- Splitting applications
- Gathering information by using BLX-SP commands

7.1 Reference Scenario

Our example is based on an actual customer running Information/Management. Before going any further, you must be aware of the main indicators of your own Information/Management system. Once you know them, you have to compare and assess your system against the example provided in this chapter before jumping to conclusions. In any case, the example will guide you by taking you throughout a step-by-step process that will help you in getting the information required to do your specific performance assessment.

In our example, the environment is characterized by the following indicators:

- 300-500 concurrent users during online time
- 1.7 million records
- 3300 new records/day
- 7000 updated records/day
- 1.7 million gets/SDIDS/day
- 200.000 puts/SDIDS/day

7.2 Collecting Installation-Dependent Definitions

Generally speaking, you may want to get definitions related to the following areas:

- MVS
- DASD
- Information/Management

In our example, MVS definitions were considered not to be relevant, so we are not including them in this redbook.

7.2.1 DASD

In terms of physical storage, our example considers the following:

- Five Cluster SDDS.

  The five clusters are installed on five volumes 3390 Model 3.

- SDIDS index and data are on separate volumes.

  The separation of index and data portions have no effect on search enqueue values.

- SDLDS placement (DASD I/O).
- Disk cache size.

7.2.1.1 VSAM Definitions

The following definitions were taken from an actual installation, but we changed all data set names to generic ones.

```
VDEF     CSECT
         SPACE 2
**********************************************************************
*     DEFINE NON-SHARED RESOURCES TO BLX-SP                          *
**********************************************************************
         BLXNSR  PLACES=40
         SPACE 2
**********************************************************************
*     DEFINE DATA SETS USING LOCAL SHARED RESOURCES (LSR)
* ----------------------------------------------------
* clusters used in different sessions
* ----------------------------------------------------
IBMPNLSW BLXDSN DSN=Wxxxxxxx.IBMPNLS.CLUSTER,                         X
               LSR=(3,3)
IBMPNLSL BLXDSN DSN=Lxxxxxxx.IBMPNLS.CLUSTER,                         X
               LSR=(5,5)
* ----------------------------------------------------
* SESSION W1 -> Production Configuration Management
* ----------------------------------------------------
SDDSW1   BLXDSN DSN=Wxxxxxxx.SDDS#01.CLUSTER,                         X
               LSR=(0,0)
SDDSW2   BLXDSN DSN=Wxxxxxxx.SDDS#02.CLUSTER,                         X
               LSR=(0,0)
SDDSW3   BLXDSN DSN=Wxxxxxxx.SDDS#03.CLUSTER,                         X
               LSR=(0,0)
SDDSW4   BLXDSN DSN=Wxxxxxxx.SDDS#04.CLUSTER,                         X
               LSR=(0,0)
SDDSW5   BLXDSN DSN=Wxxxxxxx.SDDS#05.CLUSTER,                         X
               LSR=(0,0)
SDIDSW1  BLXDSN DSN=Wxxxxxxx.SDIDS.CLUSTER,                           X
               LSR=(1,1)
SDLDSW1  BLXDSN DSN=Wxxxxxxx.SDLDS.CLUSTER,DSORG=NUM,                 X
               LSR=2
DICTW1   BLXDSN DSN=Wxxxxxxx.DICT.CLUSTER,                            X
               LSR=(4,4)
* ----------------------------------------------------
* SESSION L1 -> Education
```

```
* -------------------------------------------------
SDIDSL1  BLXDSN DSN=Lxxxxxxx.SDIDS.CLUSTER,                         X
               LSR=(7,7)
RPANELL1 BLXDSN DSN=Lxxxxxxx.PNLS.CLUSTER,                          X
               LSR=(6,6)
* -------------------------------------------------
* SESSION F1 -> Production Problem Management
* -------------------------------------------------
SDDSF1   BLXDSN DSN=Wxxxxxxx.SDDS.FMV.CLUSTER,                      X
               LSR=(0,0)
SDIDSF1  BLXDSN DSN=Wxxxxxxx.SDIDS.FMV.CLUSTER,                     X
               LSR=(1,1)
SDLDSF1  BLXDSN DSN=Wxxxxxxx.SDLDS.FMV.CLUSTER,DSORG=NUM,           X
               LSR=2
* -------------------------------------------------
* SESSION WF -> Production  IBM TRANSFER
* -------------------------------------------------
SDLDSWF  BLXDSN DSN=Wxxxxxxx.SDLDS.VER.CLUSTER,DSORG=NUM,           X
               LSR=2
*
         SPACE 2
***********************************************************************
*     GENERATE VSAM DEFINITION FOR BLX-SP                             *
***********************************************************************
         BLXGEN                           GENERATE DEFINITION
         SPACE 2
***********************************************************************
*     RESOURCE POOL ZERO (FOR PRODUCTION DATA BASE SDDS)              *
***********************************************************************
************** DATA  **************************************************
LSRD0  BLDVRP  BUFFERS=(8192(500)),     SDDS DATA BUFFERS           X
               KEYLEN=8,                POOL USED BY SDDS           X
               STRNO=30,                MAINTAIN 30 POSITIONS       X
               SHRPOOL=0,               IDENTIFY POOL               X
               TYPE=(LSR,DATA),         DATA BUFFERS ONLY           X
               RMODE31=ALL,                                         X
               MODE=24,                                             X
               MF=L
************** INDEX **************************************************
LSRI0  BLDVRP  BUFFERS=(1024(4000)),    SDDS INDEX BUFFERS AND      X
               SHRPOOL=0,               POOL IDENTIFIER             X
               TYPE=(LSR,INDEX),        INDEX BUFFERS ONLY          X
               RMODE31=ALL,                                         X
               MODE=24,                                             X
               MF=L
         SPACE 2
***********************************************************************
*     RESOURCE POOL ONE (FOR PRODUCTION DATA BASE SDIDS)              *
***********************************************************************
************** DATA  **************************************************
LSRD1  BLDVRP  BUFFERS=(16384(3000)),   SDIDS DATA BUFFERS + PAD    X
               KEYLEN=16,               POOL USED BY SDIDS          X
               STRNO=80,                MAINTAIN 80 POSITIONS       X
               SHRPOOL=1,               IDENTIFY POOL               X
               TYPE=(LSR,DATA),         DATA BUFFERS ONLY           X
               RMODE31=ALL,                                         X
               MODE=24,                                             X
```

```
               MF=L
************** INDEX **************************************************
LSRI1  BLDVRP  BUFFERS=(1024(400)),     SDIDS INDEX BUFFERS            X
               SHRPOOL=1,               POOL IDENTIFIER                X
               TYPE=(LSR,INDEX),        INDEX BUFFERS ONLY             X
               RMODE31=ALL,                                            X
               MODE=24,                                                X
               MF=L
         SPACE 2
***********************************************************************
*     RESOURCE POOL TWO (FOR PRODUCTION DATA BASE SDLDS)              *
***********************************************************************
************** DATA  **************************************************
LSR2   BLDVRP  BUFFERS=(8192(15)),      SDLDS DATA BUFFERS             X
               KEYLEN=0,                POOL USED BY SDLDS             X
               STRNO=35,                MAINTAIN 4 POSITIONS           X
               SHRPOOL=2,               IDENTIFY POOL                  X
               TYPE=(LSR),              DATA BUFFERS ONLY              X
               RMODE31=ALL,                                            X
               MODE=24,                                                X
               MF=L
         SPACE 2
***********************************************************************
*     RESOURCE POOL THREE  (PANEL DATA SETS)                          *
***********************************************************************
************** DATA  **************************************************
LSRD3  BLDVRP  BUFFERS=(4096(2000)),    PANEL DATA BUFFERS             X
               KEYLEN=10,               POOL USED BY DICT/PANELS       X
               SHRPOOL=3,               POOL IDENTIFIER                X
               STRNO=35,                COMBINED POSITIONS REQUIRED    X
               TYPE=(LSR,DATA),         INDEX AND DATA POOL            X
               RMODE31=ALL,                                            X
               MODE=24,                                                X
               MF=L
************** INDEX **************************************************
LSRI3  BLDVRP  BUFFERS=(1536(200)),     PANEL INDEX BUFFERS            X
               SHRPOOL=3,               POOL IDENTIFIER                X
               TYPE=(LSR,INDEX),        INDEX BUFFERS ONLY             X
               RMODE31=ALL,                                            X
               MODE=24,                                                X
               MF=L
         SPACE 2
***********************************************************************
*     RESOURCE POOL FOUR   (DICT  DATA SETS)                          *
***********************************************************************
************** DATA  **************************************************
LSRD4  BLDVRP  BUFFERS=(8192(100)),     DICT DATA BUFFERS              X
               KEYLEN=10,               POOL USED BY DICT              X
               STRNO=35,                MAINTAIN 5 POSITIONS           X
               SHRPOOL=4,               IDENTIFY POOL                  X
               TYPE=(LSR,DATA),         DATA BUFFERS ONLY              X
               RMODE31=ALL,                                            X
               MODE=24,                                                X
               MF=L
************** INDEX **************************************************
LSRI4  BLDVRP  BUFFERS=(1024(10)),      DICT INDEX BUFFERS             X
               SHRPOOL=4,               POOL IDENTIFIER                X
               TYPE=(LSR,INDEX),        INDEX BUFFERS ONLY             X
               RMODE31=ALL,                                            X
```

```
              MODE=24,                                              X
              MF=L
        SPACE 2
***********************************************************************
*     RESOURCE POOL FIVE (IBMPANELS)                                  *
***********************************************************************
************** DATA  **************************************************
LSRD5  BLDVRP  BUFFERS=(4096(100)),     PANEL DATA BUFFERS + PAD      X
               KEYLEN=10,               POOL USED BY SDIDS            X
               STRNO=35,                MAINTAIN 10 POSITIONS         X
               SHRPOOL=5,               IDENTIFY POOL                 X
               TYPE=(LSR,DATA),         DATA BUFFERS ONLY             X
               RMODE31=ALL,                                           X
               MODE=24,                                               X
               MF=L
************** INDEX **************************************************
LSRI5  BLDVRP  BUFFERS=(1536(12)),      PANEL INDEX BUFFERS +PAD      X
               SHRPOOL=5,               POOL IDENTIFIER               X
               TYPE=(LSR,INDEX),        INDEX BUFFERS ONLY            X
               RMODE31=ALL,                                           X
               MODE=24,                                               X
               MF=L
        SPACE 2
***********************************************************************
*     RESOURCE POOL SIX    (PANEL DATA SETS)                          *
*                             education                               *
***********************************************************************
************** DATA  **************************************************
LSRD6  BLDVRP  BUFFERS=(4096(60)),      DICT/PANEL DATA BUFFERS       X
               KEYLEN=10,               POOL USED BY DICT/PANELS      X
               SHRPOOL=6,               POOL IDENTIFIER               X
               STRNO=35,                COMBINED POSITIONS REQUIRED   X
               TYPE=(LSR,DATA),         INDEX AND DATA POOL           X
               RMODE31=ALL,                                           X
               MODE=24,                                               X
               MF=L
************** INDEX **************************************************
LSRI6  BLDVRP  BUFFERS=(1536(80)),      DICT/PANEL INDEX BUFFERS      X
               SHRPOOL=6,               POOL IDENTIFIER               X
               TYPE=(LSR,INDEX),        INDEX BUFFERS ONLY            X
               RMODE31=ALL,                                           X
               MODE=24,                                               X
               MF=L
        SPACE 2
***********************************************************************
*     RESOURCE POOL SEVEN education SDIDS
***********************************************************************
************** INDEX **************************************************
LSRD7  BLDVRP  BUFFERS=(26624(100)),    SDIDS DATA BUFFERS + PAD      X
               KEYLEN=16,               POOL USED BY SDIDS            X
               STRNO=10,                MAINTAIN 10 POSITIONS         X
               SHRPOOL=7,               IDENTIFY POOL                 X
               TYPE=(LSR,DATA),         DATA BUFFERS ONLY             X
               RMODE31=ALL,                                           X
               MODE=24,                                               X
               MF=L
************** INDEX **************************************************
LSRI7  BLDVRP  BUFFERS=(1024(100)),     SDIDS INDEX BUFFERS           X
               SHRPOOL=7,               POOL IDENTIFIER               X
```

```
         TYPE=(LSR,INDEX),        INDEX BUFFERS ONLY            X
         RMODE31=ALL,                                           X
         MODE=24,                                               X
         MF=L
   SPACE 2
***********************************************************************
   END   VDEF
```

Tip

VSAM clusters defined for the same pool must have the same CI-size or smaller or you will have errors accessing the data sets.

7.2.2 Information/Management

In Information/Management it is important to pay attention to the following parameters:

- SORTPFX
- RPANLDS
- COGENQ
- SRCHLIMIT

The above parameters are defined in the session member BLGSESW1. It is important to discuss and make an overall assessment of the implications of keeping or changing current values.

7.2.2.1 Session Member

In our example we have the definitions shown below.

```
         TITLE 'BLGSESW1 - INFO/SYSTEM SESSION PARAMETERS'
BLGSESW1 CSECT
         BLGPARMS  DICTDS=DICTDS,       NAME THE SYSTEM GLOSSARY
               RFTDS=RFTS,              NAME THE REPORT FORMAT TABLES
               PNLBCNT=250,            250 PANEL BUFFERS RESERVED
               LINECNT=100,            100 LINES
               SIDEVC=(VIO),           VIO FOR SORTINDEVICE
               SODEVC=(VIO),           VIO FOR SORTOUTDEVICE
               SWDEVC=(VIO),           VIO FOR SORTWORKDEVICE
               SORTPFX=(2500,100,100),  SORTPARAMETER
               RPANLDS=(RPANEL5,RPANEL4,RPANEL1),         READ PANELS
               DATECNV=BLGDATUM
SYSTEM   BLGCLUST NAME=5,               INFO/SYS CLUSTER
               SDDS=MGTSDDS,            NAME THE SDDS
               SDIDS=MGTSDIDS,          NAME THE SDIDS
               TRIGGER=#,               TRIGGERCHARACTER
               SDLDS=MGTSDLDS,          LOG FOR PRODUCTION
               COGENQ=15,               NEW PARM
               SRCHLIMIT=(100,100,30000)
MGTSDDS  BLGCLDSN DSN=INFOMAN.SDDS#01.CLUSTER
MGTSDIDS BLGCLDSN DSN=INFOMAN.SDIDS.CLUSTER
MGTSDLDS BLGCLDSN DSN=INFOMAN.SDLDS.CLUSTER
DICTDS   BLGCLDSN DSN=INFOMAN.DICT.CLUSTER
RFTS     BLGCLDSN DSN=INFOMAN.V6NPKMAI.RFT,FILE=RFTDD
RPANEL1  BLGCLDSN DSN=INFOMAN.IBMPNLS.CLUSTER,RDONLY=YES
```

```
RPANEL4  BLGCLDSN DSN=INFOMAN.PNLS.CLUSTER
RPANEL5  BLGCLDSN DSN=INFOMAN.BESPNLS.CLUSTER
```

There might be situations in which you can change some of the values. It is also possible that you are going to be asked to do your best without touching and changing any of the definitions.

You can find some useful hints and tips below in case you are allowed to make some changes.

Tip

- Sort Prefix Parameter
  SORTPFX=(2500,100,100)
  With these settings usesr can have search result lists of up to 2500 hits and have a maximum of 100 sorted. If possible, the value of 2500 should be lowered to a reasonable number (100-500), which means an output range of 4-20 screens.
- COGENQ=15
- Search Limit Parameter
  SRCHLIMIT=(100,100,30000)
  The application has to be analyzed to lower the value for search end (30000). Remember that search enqueue is calculated for each individual search and should be lowered as well.
- Panel Data Sets
  RPANLDS=(RPANEL5,RPANEL4,RPANEL1)
  You should reduce the amount of panel data sets to the lowest possible number, this saves time opening, searching and closing each panel data set in order to retrieve a specific panel each time.

7.3 Analyzing the Application

The first step in performance analysis and tuning has to do with data collection. In our example, we are focusing at gathering all relevant information in the two following areas:

- User dialogs and panels involved in the process
- SDIDS related information

7.3.1 Changes for Collecting Data in the BLX-SP Log

With the Search Warning Parameter (N2 of the SRCHLIMIT parameter, in the session member) you can see the user dialogs that caused excessive read operations. As of Information/Management Version 6.3 the process works as follows:

1. Message number BLG19224I is generated.
2. The message can be seen in the BLX-SP Log, the syslog, and the user is issued with a message.
3. The message contains the panel, TSP or report last used.

Since Version 6.3 you can easily change the Search Warning Parameter (N2 of SRCHLIMIT in the session parameter) then look for the output message. The user is informed when exceeding the search limit. The writing of these messages to the BLX-SP Log will not affect the system load.

Below is an example from our tests. Our search for RNID/. produced the following results:

- On the Information/Management panel the following message was shown on the Search Result List panel (BLG1TSRL):

```
BLG19224I Search Argument RNDI/. has caused an excessive number of
          database re+
```

- The following entry appeared in the SDSF LOG as well as the BLX-SP log:

```
01/10/97 10:16:04 BLG22518I Excessive number of database reads. User;
INFO1 ; SEARCH argument: RNID/. ; PANEL: BLGOE190 ; TSP/TSX: ; RFT: ;
```

For the customer reference scenario, the Search Warning Parameter was set to 100, which gave around 1000 messages during a normal day (9 hours of operation).

Be aware that message BLG22518I is a two-line message. For the technique we used, we needed just single-line messages, therefore message BLG22518I had to be modified to fit on a single line.

The following section deals with extracting the messages from the BLX-SP Log.

7.3.2 Extracting Search Warning Messages from BLX-SP Log

To have a comprehensive sample of messages it is recommended to set the SRCHLIMIT N2 parameter and have it active in a production environment for at least two hours, preferably during peak hours.

Start with a high number, for example 1000, and keep lowering it depending upon the results. It is recommended to create a specific session member for these tests.

Note: Rename your standard production session parameter while testing.

7.3.2.1 Overall Description

To have summarized information of the search warning messages you must perform the following activities:

1. Customize the BLG22518I message to make sure it fits on a single line.
2. Put the SDSF print of the BLX-SP Log to a data set.
3. Extract/compress lines with message BLG22518I using IBM DFSORT program. The DFSORT control statements can be modified according to your needs.

Notice

Message BLG22518I uses two lines because of the amount of text and data. It is recommended that you modify the message layout to fit on a single line. Another option might be to write a conversion program before you use the DFSORT utility.

7.3.2.2 How to Modify Message BLG22518I

Below are the tasks and some examples necessary to get the summary of key data.

Previous to any other tasks, you must modify the BLG22518I default message.

You can customize this message, even including national language symbols, by performing the following steps:

1. Locate the message in BLGMSGV.
2. Make the changes you want to guarantee it fits on one line.
3. Assemble and linkedit the related macro by following the recommendations suggested in BLXDMSGG.

Once you have a single line BLG22518I message, you must continue by doing the following:

1. Select **SDSF**.
2. Set Prefix to the BLX-SP name, as shown below:

```
V1R3M3 ------------------  SDSF PRIMARY OPTION MENU  ------------------
COMMAND INPUT ===> pre blx1proc                             SCROLL ===

 Type an option or command and press Enter.

    LOG       - Display the system log
    DA        - Display active users of the system
    I         - Display jobs in the JES2 input queue
    O         - Display jobs in the JES2 output queue
    H         - Display jobs in the JES2 held output queue
    ST        - Display status of jobs in the JES2 queues

    TUTOR     - Short course on SDSF (ISPF only)
    END       - Exit SDSF

    Licensed Materials - Property of IBM

    5665-488 (C) Copyright IBM Corp. 1981, 1992. All rights reserved.
    US Government Users Restricted Rights - Use, duplication or
    disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
```

3. To display BLX-SP Log enter `DA`, as shown below:

```
V1R3M3 ------------------  SDSF PRIMARY OPTION MENU  -------
COMMAND INPUT ===> da

 Type an option or command and press Enter.

    LOG       - Display the system log
    DA        - Display active users of the system
    I         - Display jobs in the JES2 input queue
    O         - Display jobs in the JES2 output queue
    H         - Display jobs in the JES2 held output queue
    ST        - Display status of jobs in the JES2 queues

    TUTOR     - Short course on SDSF (ISPF only)
    END       - Exit SDSF
```

4. Enter XD in the SDSF NP field next to the BLX-SP name to put the BLX-SP Log to a data set.

```
SDSF DA I2NO  PAGING    0  SIO     0  CPU   3%        LINE 1-1 (1)
COMMAND INPUT ===>                                           SCROLL ==
NP JOBNAME  STEPNAME PROCSTEP TYPE JNUM  C POS DP PGN REAL PAGING   SIO
xd BLX1PROC BLX1PROC BLXSPCAS STC  21113   NS  57   5 7924   0.00   0.0
```

The SDSF Open Print Data Set panel is now displayed. Fill in the the information needed. Use the panel shown below as an example.

```


                            SDSF Open Print Data Set
COMMAND INPUT ===>                                           SCROLL ==


Data set name  ===> 'B5M1GRA.SEA'
Member to use  ===>
Disposition    ===> OLD        (OLD, NEW, SHR, MOD)

If the data set is to be created, specify the following:

Management class     ===>           (Blank for default management class
Storage class        ===>           (Blank for default storage class)
  Volume serial      ===>           (Blank for authorized default volum
Data class           ===>           (Blank for default data class
  Space units        ===> BLKS      (BLKS, TRKS, CYLS, KB, or MB)
  Primary quantity   ===> 500       (In above units)
  Secondary quantity ===> 500       (In above units)
  Directory blocks   ===>           (Zero for sequential data set)
  Record format      ===> FB
  Record length      ===> 240
  Block size         ===> 3120
```

5. To activate the entries and print the BLX-SP Log press Enter.

6. The SDSF panel nnn LINES PRINTED comes up now.

```
SDSF DA I2NO  PAGING    0  SIO    71  CPU   7%      7,789 LINES PRINTED
COMMAND INPUT ===>                                           SCROLL ===
NP JOBNAME  STEPNAME PROCSTEP TYPE JNUM  C POS DP PGN REAL PAGING   SIO
   BLX1PROC BLX1PROC BLXSPCAS STC  21113   NS  57   5 8263   0.00   9.91
```

7. To close the print data set enter XC in the SDSF NP field, or close print on the command line.

```
SDSF DA I2NO  PAGING    0  SIO    71  CPU   7%      7,789 LINES PRINTED
COMMAND INPUT ===>                                           SCROLL ===
NP JOBNAME  STEPNAME PROCSTEP TYPE JNUM  C POS DP PGN REAL PAGING   SIO
XC BLX1PROC BLX1PROC BLXSPCAS STC  21113   NS  57   5 8263   0.00   9.91
```

8. The SDSF message PRINT CLOSED is generated.

```
SDSF DA I2NO  PAGING    0  SIO    82  CPU   3%         PRINT CLOSED  155
COMMAND INPUT ===>                                          SCROLL ===
NP JOBNAME  STEPNAME PROCSTEP TYPE JNUM  C POS DP PGN REAL PAGING   SIO
   BLX1PROC BLX1PROC BLXSPCAS STC  21113   NS  51   5 8498   0.00  16.17
```

9. Submit the extract job.

   This job might need to be modified according to your installation needs and the layout of your own BLG22518I message. We used the INREC FIELD definitions shown below to select the information we wanted. You may want to have a look at the IBM DFSORT brochure or documentation in case you want to make changes to the definitions included below.

```
//B5M1GRAS JOB 7051000S,'GRAEF',
//             NOTIFY=B5M1GRA,
//             CLASS=A,MSGCLASS=T
//SORT   EXEC PGM=SORT
//SYSPRINT DD SYSOUT=*
//SYSOUT   DD SYSOUT=*
//SORTLIB  DD DISP=SHR,DSN=SYS1.SORTLIB
//SORTIN   DD DISP=SHR,DSN=B5M1GRA.SEA
//SORTOUT  DD DSN=B5M1GRA.SEARCH.SORT1910,
//            DISP=(,CATLG),UNIT=SYSDA,
//            SPACE=(CYL,(1,1)),
//            DCB=(DSORG=PS,RECFM=FB,LRECL=51,BLKSIZE=10200)
//SORTWK01 DD UNIT=SYSDA,SPACE=(CYL,1)
//SORTWK02 DD UNIT=SYSDA,SPACE=(CYL,1)
//SORTWK03 DD UNIT=SYSDA,SPACE=(CYL,1)
//SYSIN    DD *
 INREC FIELDS=(5,2,10,1,11,5,19,1,20,10,46,80)
 INCLUDE COND=(20,9,CH,EQ,C'BLG22518I')
 SORT FIELDS=(20,5,CH,A)
 END
```

7.3.2.3 Expected Results of the Extract Job

The output was sorted for p-words (col. 20-23) in such a way that:

- Column 1-2 is the recorded day.
- Column 4-8 is the message time.
- Column 10-19 is the message number.
- Column 20 is the p-word, the search argument used followed by the panel, TSP or RFT name.

You can find below the result of our extract job sample.

```
Cols 01-02 = Date
     03-07 = Time
     09-17 = Message number
     19-   = Search argument, panel- ,RFT- or TSPname


19 11:39 BLG22518I ABEZ/. from panel MC66RFTN, TSP BLGTENDS or RFT
19 11:46 BLG22518I ABEZ/. from panel , TSP BLGTENDS or RFT MDKFPCK
19 10:33 BLG22518I ABEZ/. from panel MC66RFTN, TSP BLGTENDS or RFT
19 08:47 BLG22518I DATX/. from panel , TSP AUFTA201 or RFT  has caused a
19 08:47 BLG22518I DATX/. from panel , TSP AUFTA201 or RFT  has caused a
19 15:09 BLG22518I DSTE/. from panel , TSP BTNTENTR or RFT  has caused a
19 15:09 BLG22518I DSTE/. from panel , TSP BTNTENTR or RFT  has caused a
19 15:10 BLG22518I DSTE/. from panel , TSP BTNTENTR or RFT ........ has
19 15:10 BLG22518I DSTE/. from panel , TSP BTNTENTR or RFT ........ has
19 16:31 BLG22518I DTED/GKV. from panel , TSP USRTA104 or RFT  has cause
19 07:52 BLG22518I DTED/RE. from panel , TSP BLGTENDS or RFT  has caused
19 16:31 BLG22518I DTED/GKV. from panel , TSP USRTA104 or RFT  has cause
19 07:52 BLG22518I DTED/RE. from panel , TSP BLGTENDS or RFT  has caused
19 12:36 BLG22518I GEZI/1. from panel MC623401, TSP BLGTENDS or RFT  has
19 12:24 BLG22518I GEZI/1. from panel MC623401, TSP BLGTENDS or RFT  has
19 12:23 BLG22518I GEZI/2. from panel MC623401, TSP BLGTENDS or RFT  has
```

Figure 10. Example of Output of the Search Limit Warning Message

7.3.2.4 Considerations for the Usage of P-Words and S-Words

To use output shown in Figure 10 as an example, you can see that p-word ABEZ was used for a range search in report MDKFPCK. Also that p-word DATX was used for a range search with TSP AUFTA201.

If you have different environments running with the same BLX-SP, all information is in the same LOG. The following points have to be taken into consideration:

1. How many read operations do I perform with that range search?

   There are more choices available to come to a result:

   - Look at the Glossary and count the entries for that p-word. Each entry reflects one read operation. This is not preferable for hundreds of entries.
   - Run the dialog as a single user and enter the BLX-SP Query command before and after the dialog:

     `/F procname,QUERY,FILE=xxxx,TYPE=VSAM`

     Note: The / only has to be used when issuing the command via SDSF. To find out the xxxx value, you can use the Query command, which will give you the list of VSAM definitions (check the SDSF syslog or BLX-SP log):

     `/F procname,QUERY`

     Note: Substitute your BLX-SP name for procname.

     When you use this procedure you have to subtract the read operations after doing the search again.

   - Use the BLGUT21 procedure for stats on the SDIDS. Run BLGUT21 in analyze mode then change the previously used DFSORT batch job to extract the p-word you are looking for and count the lines of output. Each line represents one read operation to the SDIDS. To search for the presence of a field you should never use the p-word, which causes massive read operations thus slowing down response. In these cases you should use the s-word, which results in one read operation only.

2. If you decide to use the adjacent s-word, you have to consider the following:

   - If the s-word word is still in the database, then just use it in RFT or TSP. (See *Terminal Simulator Guide and Reference*, SC34-4609 for how to use s-words in a TSP.) When using S-words in TSP you have to insert blanks before and after.
   - If the s-word was removed during the p-word cognize only process, then you have to decide what the greatest benefit for the application is. Adding the s-word again needs one read and one write operation during create or update process for one record. You have to put the amount of extra read operations against this.

7.4 Analyzing SDIDS for P-Words

After obtaining all relevant data from search warning messages, you may want to know how many p-words are used in SDIDS data sets. You have two choices to do this:

- Search the glossary
- Use the BLGUT21 utility in analyze mode

Of the above options, we recommend the second one. Even though the first option is a valid one, we deem it more cumbersome and time consuming.

7.4.1 Procedure to Extract Data from SDIDS

The procedure described in the following sections is only related to the use of the BLGUT21 utility to collect the proper information from a SDIDS data set:

1. Run BLGUT21 in analyze mode.
2. Modify/run IBM DFSORT.

7.4.1.1 Detailed Description to Collect SDIDS Information

Below are the steps, with some examples, to extract the appropriate data relating to p-words:

1. Run BLGUT21 in analyze mode.

   To run the Information/Management utility in analyze mode you need 250 cylinders of 3390 Model 3 to keep 3.5 million records.

```
//BLGUT21  JOB ,JOBCARD
//UT21    EXEC PGM=BLGUT21,REGION=2048K
//STEPLIB  DD DISP=SHR,DSN=SYS1.SBLMMOD1
//SYSPRINT DD SYSOUT=*
//BLGSDIDS DD DISP=SHR,DSN=BLM.V6R3M0.SDIDS
//ANALYZE  DD DSN=B5M1GRA.ANALYZE,
//            DISP=(,CATLG),UNIT=SYSDA,
//            SPACE=(CYL,(250,1)),
//            DCB=(DSORG=PS,RECFM=FB,LRECL=51,BLKSIZE=10200)
//SYSIN DD DUMMY
//
```

2. Extract all p-words with RNSY/.

With the job shown below you can select all p-word entries for the p-word RNSY. Each line represents one data value for this p-word. The definitions are made with the INCLUDE and SORT FIELD statements.

```
  Sortparameter: select from column 11 length 5
                 value RNSY/
                 sort column 11 length 5

//DFSORT   JOB ,JOBCARD
//SORT   EXEC PGM=SORT
//SYSPRINT DD SYSOUT=*
//SYSOUT   DD SYSOUT=*
//SORTLIB  DD DISP=SHR,DSN=SYS1.SORTLIB
//SORTWK01 DD UNIT=SYSDA,SPACE=(CYL,1)
//SORTWK02 DD UNIT=SYSDA,SPACE=(CYL,1)
//SORTWK03 DD UNIT=SYSDA,SPACE=(CYL,1)
//SORTIN   DD DISP=SHR,DSN=B5M1GRA.ANALYZE
//SORTOUT  DD DSN=B5M1GRA.ANALYZE.SORT,
//            DISP=(,CATLG),UNIT=SYSDA,
//            SPACE=(CYL,(1,1)),
//            DCB=(DSORG=PS,RECFM=FB,LRECL=51,BLKSIZE=10200)
//SYSIN DD *
 INCLUDE COND=(11,5,CH,EQ,C'RNSY/')
 SORT FIELDS=(11,10,CH,A)
//
```

3. Gather the p-word-related information.

   The output looks like the example below:

```
000001 00000028  RNSY/00000000
000002 00000028  RNSY/01523854
000003 00000028  RNSY/01916721
000004 00000028  RNSY/01923949
000005 00000028  RNSY/01524212
000006 00000028  RNSY/01891841
000007 00000028  RNSY/01880143
000008 00000028  RNSY/01896943
000009 00000028  RNSY/01896964
000010 00000028  RNSY/01896973
000011 00000028  RNSY/01896984
000012 00000028  RNSY/01897012
000013 00000028  RNSY/01897027
000014 00000028  RNSY/01897075
000015 00000028  RNSY/01897079
000016 00000028  RNSY/01911265
```

7.5 Analyzing Report Data Sets for Range Searches

To check whether your reports are using range searches or not you simply can use this procedure:

1. Under ISPF, select **Utilities** (Selection 3).
2. In the Utility Selection menu, select **Search For** (Selection 14).
3. The following screen is displayed:

```
----------------------------- SEARCH-FOR UTILITY --------------- STRINGS
COMMAND ===>

SEARCH STRING    ===> 'SEARCH ARGUMENT('

MULTIPLE STRINGS ===> NO   (Yes to specify additional search strings)

ISPF LIBRARY:
   PROJECT ===>
   GROUP   ===>          ===>          ===>          ===>
   TYPE    ===>
   MEMBER  ===>                  (Blank or pattern for member selection
                                  '*' for all members)
OTHER PARTITIONED OR SEQUENTIAL DATA SET:
   DATA SET NAME  ===> 'W5951PDS.V6NPKMAI.RFT(*)'
   VOLUME SERIAL  ===>           (If not cataloged)

DATA SET PASSWORD ===>           (If password protected)

LISTING DSNAME    ===> SRCHFOR.LISTN
MODE              ===> F         (F - foreground, B - batch)
MIXED MODE        ===> YES       (Yes or No)
```

4. Fill in the appropriate values and press Enter.
5. The following is a sample output.

   The report member name and the search argument are shown. To browse through the output you can simply use the find command to search for wild search values (for example RNID/.).

```
This is the output data set with the search arguments listed :

 BLMZZ23              ----------- STRING(S) FOUND ------------------

    598  SEARCH ARGUMENT(!SOBO1 ¬ !SOB5A ¬ !SOB9C ¬ STAC/CLOSED)

BLMZZ23 is the report member name
```

Figure 11. Output Search For Utility

7.6 Splitting Applications

In complex and heavily loaded Information/Management installations you may consider improving performance by designing a new system in which different applications have separate access to VSAM data sets. You can see in Figure 12 on page 72 an example of this approach, taken from a real customer scenario.

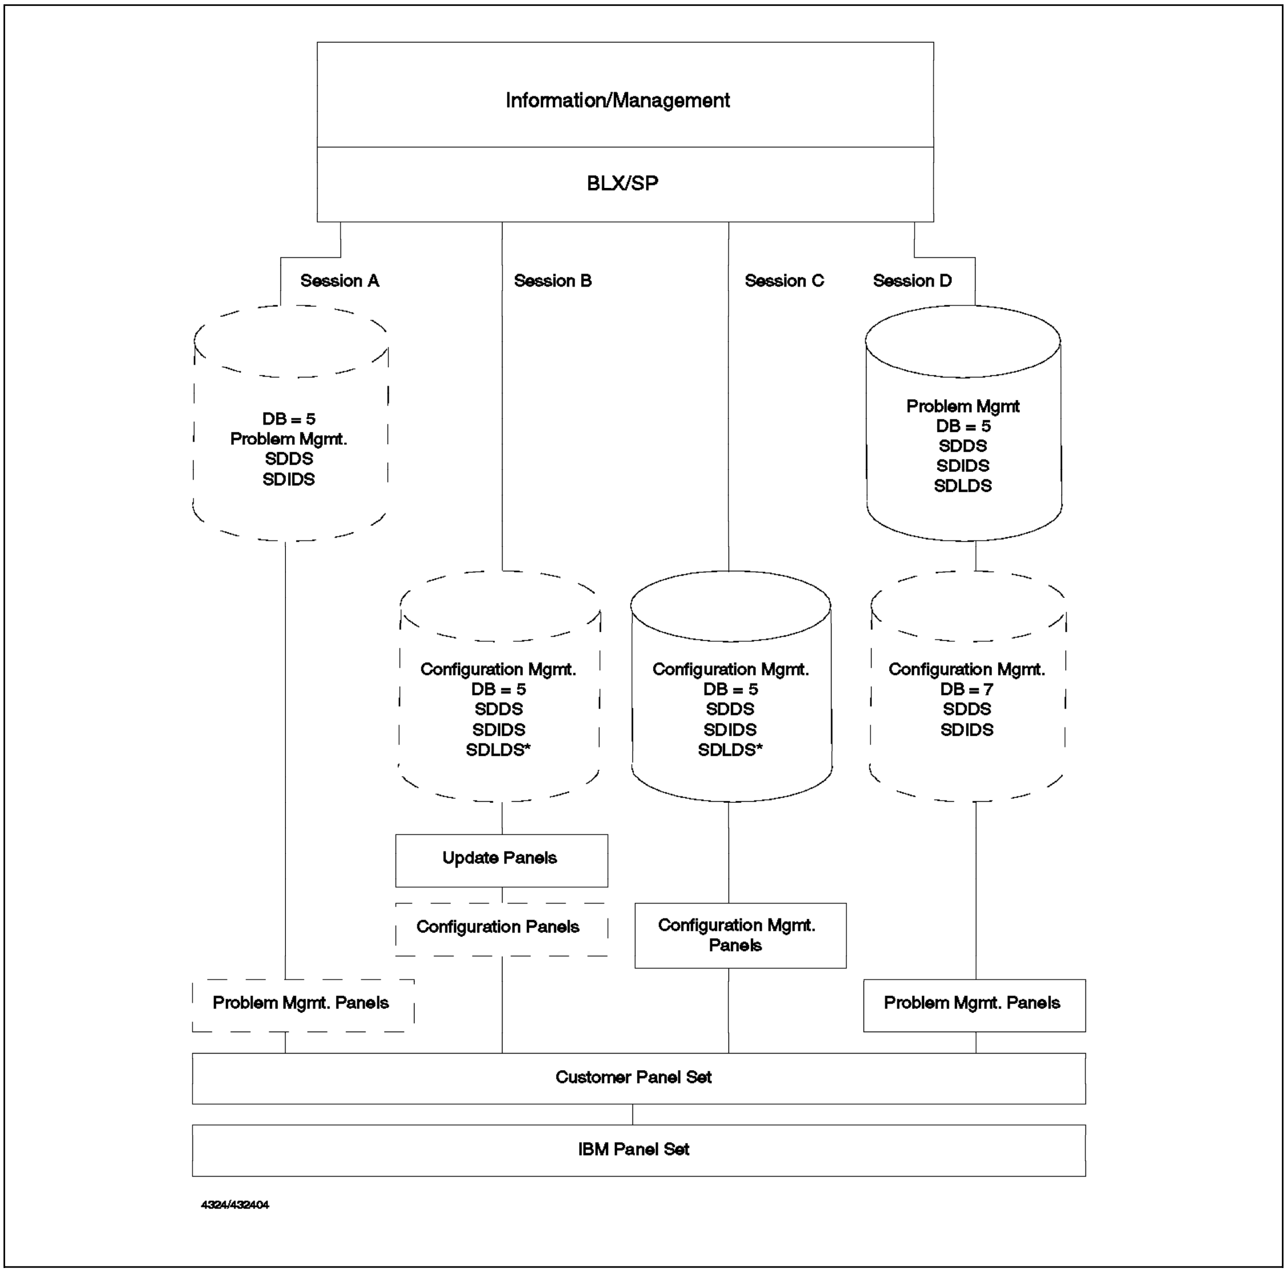

Figure 12. Example from a Customer Study

With Information/Management you can have a session with read/write access to a database and read-only access to other databases. Because the main bottleneck for Information/Management is the access to SDIDS the only way to enhance performance is by creating a separate session with its own set of clusters (SDDS,SDIDS and SDLDS). The panel data set(s) can be the same. You can define the same cluster in different sessions with different access characteristics for example, in session A as a read/write cluster and in session B as read only.

If you have two clusters defined in one session (for example DB=5 and DB=7), then you have two choices to access the different databases:

1. Temporarily change your profile
2. Set up your search argument like this:

```
/se =7 stac/open
```

7.6.1 Customer Case Reference

In the actual customer environment, after taking into account different alternatives, the final decision was to separate problem management from the other applications (which represented 25% of the load).

The solution had to be implemented based upon the following customer requests being fulfilled:

1. User should not see any difference.
2. Privilege classes are updated in the Configuration/Management Session only and have to be transferred with batch jobs to the Problem/Management Database. To make the usage as easy as possible the class records in both sessions have the same authority. The access has to be controlled by TSPs.
3. Configuration/Management read has to be possible from the Problem/Management Session.
4. Configuration/Management write is only possible from the Configuration/Management Session.
5. Problem/Management record IDs have to start with F. (The reason for this request is to possibly combine the problem and configuration records in one database.)
6. Problem records must be transferred to the new database and have the same RNID as in the old database.

7.6.2 Description of a Splitting an Application Example

Figure 12 on page 72 depicts a solution taken from an actual customer, in which:

- Dotted clusters are logical, which means the same clusters are defined in different sessions but physically they are the same. In terms of access, dotted clusters have read only access, while whole clusters have read/write access.
- Clusters defined as database 5 are read/write databases, whereas clusters defined as database 7 are read only. It is no problem to define the same cluster in one session as database 5 and in another session as database 7.
- The SDLDS that is marked as * is a temporary data set, which is used for class record transfer only using BLGUT4 and BLGUT3.
- Session A
  - This session is used to transfer Problem/Management records during installation to the Configuration/Management database. This session runs without an SDLDS data set, so that if there are any problems encountered during the transfer process it is easier to rerun the complete job. For this session there′s batch program access only.
- Session B
  - This session is used to unload the records that have to be transferred from the Configuration/Management database to the Problem/Management database. For this session a special log data set is used that is not under control of the normal backup procedure. This log data set is used as a temporary data set only. TSPs are used to

perform dummy updates so that the records are saved to the SDLDS. This was done for the first time transfer then the daily updates were used for the changed class records. The Update Panel data set allows updates for Problem/Management records in the Configuration/Management database.

- Session C
  - This is the normal production session for Configuration/Management.
  - Configuration/Management database is defined for read/write access DB=5.
- Session D
  - This is the normal production session for Problem/Management.
  - The Configuration/Management database is defined for read only access DB=7.
  - The Problem/Management database is defined for read/write access DB=5.
  - The old Problem/Management records (transferred) have numerical RNIDs.
  - The new Problem/Management records have numerical RNIDs starting with F.

7.6.3 Actions to Perform

You must make some changes to the following sessions:

- Configuration/Management session
- Problem/Management session

7.6.3.1 Changes to the Configuration/Management Session

In this subject, you must remove all selections that relate to Problem/Management. Another possibility is disabling the authority for Problem/Management in the class record.

7.6.3.2 Changes to the Problem/Management Session

To take advantage of the new option, you must make a temporary profile change (to database 7, instead of 5) to select the correct database before each access.

7.7 Gathering Performance Data By Using BLX-SP Commands

Information/Management doesn′t come with a comprehensive set of tools for performance measurement, accounting, or other operations and planning-related areas. However, you may want to use some techniques based on operator commands to create your own historical database from which you could get the information necessary for performance tuning and capacity planning purposes.

In this section you can find guidance to help you get performance information from BLX-SP commands. As the command output only gives you the data at the moment the command was issued, that means that you are using a technique based on samples, therefore, you must make sure you collect the information periodically, let′s say, every half an hour.

For a complete command description have a look at *Operation and Maintenance Reference*, SC34-4480. In the following sections we describe the main important tasks for performance issues only. Notice that the counters in BLX-SP output are cumulative, so you have to calculate the difference between two consecutive command runs. Another possibility is to set up the commands with any automation tool at a specific time and count/calculate the difference. The data set we are looking at is the SDIDS.

The maximum command length you can enter in SDSF (online or by batchjob) is 42 characters. So you should use the file name for a cluster instead of the data set name. To find the file name for a specific cluster, for example the SDIDS, run the query command without any parameters:

```
/F procname,QUERY
```

You may want to use a similar procedure to what it is described in 7.3.2.2, "How to Modify Message BLG22518I" on page 65 to extract data from the log.

7.7.1 User Count

You can find out the number of users logged on to Information/Management by issuing the following command:

```
/F procname,QUERY,FILE=SDIDSW1,TYPE=STATUS
```

In the example below, you can see that the number of users is 239.

```
07:00:01 QUERY,FILE=SDIDSW1,TYPE=STATUS
07:00:01 BLX03220I STATUS
07:00:01 BLX03209I File: SDIDSW1          Data set: W2951WDB.SDIDS.CLUSTER
07:00:01 BLX03221I      USERS     ACTIVE    RESPOOL     PHMAXW      PHMAXU
07:00:01 BLX03222I        239          6          1          0        15
```

7.7.2 Type=IO

You can find out the I/O activity of the SDIDSW1 data set by issuing the following command:

```
/F procname,QUERY,FILE=SDIDSW1,TYPE=IO
```

```
07:00:01 QUERY,FILE=SDIDSW1,TYPE=IO
07:00:01 BLX03215I IO statistics
07:00:01 BLX03209I File: SDIDSW1          Data set: W2951WDB.SDIDS.CLUSTER
07:00:01 BLX03216I      OPENS     CLOSES       GETS       PUTS
07:00:01 BLX03217I       7325       7324     612264       1933
07:00:01 BLX03218I     ERASES    RETRIES     POINTS    ENDREQS
07:00:01 BLX03219I         14          0      19154      33566
```

In the example, the output gives you information about:

- Open/closes

The IOs to VSAM are made from BLX-SP and not from the user. So these values are logical opens and closes which you don′t have to worry about.

- Gets

  These are the total reads to that cluster. As already explained you have to run this command twice during a specific period and subtract the values from each other to have the read operations during a time period.

- Puts

  These are the total write operations to that cluster.

7.7.3 Type=VSAM

You can find out the VSAM activity of the SDIDSW1 data set by issuing the following command:

```
/F procname,QUERY,FILE=SDIDSW1,TYPE=VSAM
```

```
07:00:01 QUERY,FILE=SDIDSW1,TYPE=VSAM
07:00:01 BLX03223I VSAM statistics:
07:00:01 BLX03228I Shared Resource Pool: 1.
07:00:01 BLX03229I      PHACT     PHMAXW     PHMAXU
07:00:01 BLX03230I          1          0          3
07:00:01 BLX03209I File: SDIDSW1          Data set: W2951WDB.SDIDS.CLUSTER
07:00:01 BLX03224I      HURBA   CASPLITS    EXTENTS    GETSWIO
07:00:01 BLX03225I  412876800          1          1      65632
07:00:01 BLX03226I      HARBA   CISPLITS     STRMAX   GETSNOIO
07:00:01 BLX03227I  442368000         20          3      17187
10:00:01 QUERY,FILE=SDIDSW1
```

In the example, the output gives the following information about:

- GETSWIO

  These are the read operations to the disk (or if defined disk cache).

- GETSNOIO

  These are the read operations that are satisfied by the LSR buffer. Remember that if GETSWIO is higher than GETSNOIO it could be due to:

  – The data is not in LSR (for example a just updated record).

  – A spanned record has to be read.

- CASPLITS

  Avoid CASPLITS if possible as they severely degrade performance.

  Note: The only way to reorganize this data set in a spanned environment is to delete/define the cluster and run BLGUT1 afterwards.

- CISPLITS

  CISPLITS are not as critical as CASPLITS but you should try to reduce them as well.

Tip

In the reference customer installation the hit ratio was 26 % (comparing GETSWIO and GETSNOIO). If this figure remains constant, the LSR buffers should be lowered.

7.7.4 Commands TYPE=IO and TYPE=VSAM Three Hours Later

Below you can see the output of the QUERY command, which was issued three hours later (at 10 a.m.):

```
/F procname,QUERY,FILE=SDIDSW1,TYPE=IO
```

```
10:00:01 QUERY,FILE=SDIDSW1,TYPE=IO
10:00:01 BLX03215I IO statistics
10:00:01 BLX03209I File: SDIDSW1           Data set: W2951WDB.SDIDS.CLUSTER
10:00:01 BLX03216I      OPENS     CLOSES       GETS       PUTS
10:00:01 BLX03217I      30473      30467     935511      60545
10:00:01 BLX03218I     ERASES    RETRIES     POINTS    ENDREQS
10:00:01 BLX03219I        229          0      92702     166569
```

Similarly, we did the same for the VSAM command:

```
/F procname,QUERY,FILE=SDIDSW1,TYPE=VSAM
```

The command response is shown below:

```
10:00:01 QUERY,FILE=SDIDSW1,TYPE=VSAM
10:00:01 BLX03223I VSAM statistics:
10:00:01 BLX03228I Shared Resource Pool: 1.
10:00:01 BLX03229I      PHACT     PHMAXW     PHMAXU
10:00:01 BLX03230I          6          0         15
10:00:01 BLX03209I File: SDIDSW1           Data set: W2951WDB.SDIDS.CLUSTER
10:00:01 BLX03224I      HURBA   CASPLITS    EXTENTS    GETSWIO
10:00:01 BLX03225I  412876800          1          1     660903
10:00:01 BLX03226I      HARBA   CISPLITS     STRMAX   GETSNOIO
10:00:01 BLX03227I  442368000         23          9     122182
```

7.7.5 Analysis of Results After 3 Hours

As mentioned before all counters are cumulative, so we had to calculate the difference to get the actual activity, which is shown below:

- Total GETS (reads): 323247
- Total PUTS (writes): 58612
- Total GETSWIO: 595271
- Total GETSNOIO: 104995

At a first glance the figures Total GETS, Total GETSWIO and Total GETSNOIO don′t match. The explanation for these different figures is the following: GETSWIO and GETSNOIO reflect physical VSAM record accesses, GETS and PUTS reflect logical record accesses. Therefore they don′t take into account anything about spanned records, etc.

Conclusion

The customer was suffering real performance problems. Most user activities were record create and update. The read-to-write ratio is 5.5. The amount of read operations is too high. The conclusion is that the application logic had to be analyzed and the code re-written to reduce unnecessary read operations.

7.8 Using TASID - Examples

TASID is a very simple but powerful tool that is available within IBM as TASID PACKAGE from the MVSTOOLS repository. It can be obtained from your IBM representative. Consider getting a copy if you are going to do some serious performance analysis.

The best way of make people aware of the use of tools is by using examples to illustrate a few actions and results. All data was taken from a running system that has TASID installed.

To request performance information you have to type `TSO TASID` on any TSO screen and press Enter. The result will be the TASID main screen as shown in Figure 13.

```
CPU=27%                         TASID option menu
Option  ===>

   Select one of the following options:                   Version 5.04
       1 - Address space list        5 - Miscellaneous displays
       2 - System ENQ contention     6 - Current data set allocations
       3 - Total system ENQ status   7 - Storage View Facility
       4 - Initiator Status List     8 - Snapshot
                                                                   More:
  +--------------------------------------+------------------------------+
  !   Current time  15:36 on 97/01/07    !   TSO users         250     !
  !   Last IPL time 09:52 on 97/01/07    !   Started tasks      40     !
  !   System   SP4.3.0  JES version JES2   !   Jobs                1     !
  !   SMF ID   I2N0     JES level   4.3.0  !   System addrs       17     !
  !   User ID  N28PA02  RACF level  1.9.2  !   Free initiators    33     !
  !   Node     A8PNJES2 TSO version 2.04.0 !------------------------------!
  !   VTAM Adr A8SAXG04 VTAM Level  4.2.   !   Total             341     !
  !   Proc     $IRIS    DFSMS level 1.1.10 !------------------------------!
  !   Region   4M                          !   CPU utilization 27%       !
  !   RACF Grp VOR1     DSF  level  1.16.0 !   CPU 9021 (running 1 CPUs)  !
  !                     ISPF Level  3.5.0  !   Real Storage    1048576K   !
  !                     HSM Level   1.1.0  !   ENQ Contention  1          !
  !--------------------------------------+-------------------------------!
  !   MVS Information: MVS/ESA SP4.3.0 HBB4430                          !
  !   JES Information: SP 4.3.0                                         !
  !---------------------------------------------------------------------!
  !   This system keeps a history of 5 passwords.                       !
  !   Automatic revocation after 3 invalid logon attempts.              !
  ----------------------------------------------------------------------!
```

Figure 13. TASID Main Screen

From the TASID main screen you can get two important performance indicators:

- Logged on TSO users = 250 (users in that LPAR)
- CPU/LPAR utilization = 27%

The TASID main screen allows you to request more detailed information by means of options. For Information/Management purposes it is worth looking at option 3 (total system ENQ status) and option 2 (system ENQ contention).

In our test we selected option number 3 by typing 3 and pressing Enter. The result is shown in Figure 14.

```
CPU=49%                        System ENQ Status              Line 1 of
Command ===>                                                  Scroll ===

  Major name prefix ===> ________   (SYSDSN, SPFEDIT, etc)
  Minor name prefix ===> ____________________________________________ (d
  Address id prefix ===> ________   (Job name, User id, etc)


  Major      Minor                                               Job

  .H15     | SUBSYSTEM                                        | RDS
  $CACHENV | .ô.ø.pÿ.  ′00ED3E800E97DF38′X                    | OMCMS
  AFOPERAT | OG1                                              | OSM1
  ARCGPA   | ARCCAT                                           | DFHSM
  BETASFM  | B44                                              | BETA44
  BETASFM  | B92                                              | BETA92
  BLGRNID  | W2951WDB.SDDS.FMV.CLUSTER(F0502606)              | XK2I012
  BLGRNID  | W2951WDB.SDDS.FMV.CLUSTER(F0502680)              | N3IKZ21
  BLGRNID  | W2951WDB.SDDS.FMV.CLUSTER(F0502795)              | XC20261
  BLGRNID  | W2951WDB.SDDS.FMV.CLUSTER(F0502798)              | XEW0061
  BLGRNID  | W2951WDB.SDDS.FMV.CLUSTER(F0502814)              | X9E0172
  BLGRNID  | W2951WDB.SDDS.FMV.CLUSTER(F0502819)              | NOFOY20
  BLGRNID  | W2951WDB.SDDS.FMV.CLUSTER(F0502863)              | XG20170
  BLGRNID  | W2951WDB.SDDS.FMV.CLUSTER(F0502866)              | XSU0019
  BLGRNID  | W2951WDB.SDDS.FMV.CLUSTER(F0502868)              | XS20929
  BLGRNID  | W2951WDB.SDDS.FMV.CLUSTER(F0502869)              | XVE0540
  BLGSEQN  | W2951WDB.SDDS.FMV.CLUSTER(00006FCB)              | N1M4B08
  BLGSEQN  | W2951WDB.SDDS.FMV.CLUSTER(00007E05)              | XMR0012
```

Figure 14. TASID Screen for Option 3

The tool provides detailed data showing ENQ status for clusters, RNIDs and users. For example, RNID F0502606 (from BLGRNID group) is held by user XK2I012.

If instead of option 3 you decide to select option 2 from the main screen, the result would be a screen as the one shown in Figure 15 on page 80.

Be aware that ENQ status on SDIDS may be a normal situation in a running system. However, the queue should not exceed a reasonable number, let's say 10. If there is no queue condition, you will get a screen with the message `No ENQ Condition`.

```
CPU=42%                     System ENQ contention                Lin
Command ===>                                                 Scroll ==

      Data set name                                  Owner      Request
     ┌──────────────────────────────────────────────┬──────────┬──────────┐
     │ W2951WDB.SDIDS.CLUSTER                        │ N2P1Y11  │ XA20083  │
     │                                               │          │ XC20249  │
     │                                               │          │ XA20535  │
     │                                               │          │ NKLKB01  │
     └──────────────────────────────────────────────┴──────────┴──────────┘
```

Figure 15. TASID Screen for Option 2

It is mandatory to use 3270 color displays or emulators as some key information is provided by means of colors. Figure 15 presented some information as follows:

- User N2P1Y11 is shown in red, meaning write ENQ.
- XA20083 is shown in white, meaning read request.
- XA20535 is shown in read, meaning write request.

The example tells you that a write operation holds an exclusive ENQ to the SDIDS.

Appendix A. Searching and Sorting Flash

The information in this appendix has been previously published as an ITSO flash and is reproduced here because of its relevance.

A.1 Information/Management Search/Report Performance

This flash discusses the performance implications of certain types of Information/Management search and report transactions. It provides advice on recognizing whether your Information/Management performance is affected by such transactions and how to avoid or minimize the impact.

A.1.1 How Search/Report Performance Can Be Impaired

Search and report transactions each have several steps:

- Identify the records in the SDDS which match the search criteria.
- Provide a way of retrieving the records in the proper order, if the results are to be sorted.
- Read and display or print the required records.

Performance of search and report transactions is affected by all three of these steps. Usually, Information/Management can handle all of these steps very efficiently. Several techniques are used to try to process the transactions in the most efficient way:

- Search criteria resolution for either search or report transactions is done by using the SDIDS records as a very efficient index into the records in the SDDS. For most search criteria, no matter how many records or bytes are in your SDDS, Information/Management can determine which records you were searching for by reading just a few SDIDS records.
- For search transactions, Information/Management also attempts to choose the most efficient sort technique, when optional sorting is specified. If your parameters and sort criteria are reasonable, the sorting of very large Search Results Lists (SRLs) can be done by reading only a few SDIDS records rather than all the SDDS records.

There are special cases where these two activities can become very expensive operations and performance can be affected. These conditions occur when massive numbers of SDIDS records may be required to satisfy the search criteria of a search or report or to complete a sort of a large SRL. In such cases the response times for the search or report transaction will be significantly longer than a ″normal″ search or report, which does not require these massive numbers of I/Os. Further, since the SDIDS is frequently a key resource, the resulting enqueue contention for the SDIDS can have a significant impact on the performance of other Information/Management users.

A.1.2 What Does This Mean to Versions 4, 5, and 6?

The way these items work in Information/Management has not really changed, but the processor (CPU) costs of doing each VSAM I/O are higher starting with Version 5. While most workloads as a whole will only show modest increases in absolute CPU cost, if search or sort activities require massive SDIDS reads, then significant increases in response times and CPU use can occur.

The performance improvements implemented in Version 5 (and continuing into Version 6) were accomplished by avoiding or speeding up VSAM I/O with smarter buffering through a central address space. The reductions in I/O times can result in dramatically shorter response times in Versions 5 and 6. Version 5 maximum throughput can be double or triple that of Version 4, and even greater increases in throughput are possible with Version 6.

For most transactions, the extra CPU costs of each VSAM I/O are not a serious problem, maybe completely unnoticeable. However, where search criteria or SRL sorting causes massive numbers of SDIDS reads, the extra cost can be high. Although this sounds like it would stress I/O resources, the records are buffered so well that CPU resources (instead of I/O) can become stressed during such processing.

The effect can be longer response times for some search or report transactions, which cause massive SDIDS reads. The increased enqueue contention for the SDIDS can also delay other users who need the SDIDS. For transactions that can cause massive numbers of SDIDS reads, Versions 5 and 6 can actually result in worse performance than in earlier versions. Fortunately, massive SDIDS reads can usually be avoided, and the dramatic Version 5 and 6 performance improvements do not need to be compromised.

Whether you are currently running or plan to move to Version 5 or Version 6, this flash can help you recognize and avoid performance costs from some special types of search or sort activities.

A.1.3 What Can Cause Massive Numbers of SDIDS Reads?

There are three primary causes for this situation:

- **Big-range searches:** Whenever the search criteria of a search or report transaction specifies a range of values on which to search, all the SDIDS records must be read. For example, if you entered the search as follows:

  ```
  SE DATE/92/11/01 - 92/12/31
  ```

  , then to identify the hits, 62 SDIDS records would have to be read. That's one for the master bit list, and one for each of the different days in the specified date range. Note that this is *not* a big range search, and response time would be very fast. However, suppose your search range was `SE RNID/R0000001 - R01000000` and you actually had a million records in your database. Satisfying that search would require that one million SDIDS records be read, which is certainly a massive read situation.

- **Big period or asterisk searches:** Using periods (.) or asterisks (*) as wildcard characters in search criteria is not usually a problem, but it can become one if many thousands of SDIDS records must be read. All SDIDS records defined by the wildcard range must be read. For example, `SE DATE/92.` would result in a maximum of 366 SDIDS records being read, which, again, is not a problem. However, `SE RNID/.` would result in one SDIDS record read for each record in your entire SDDS: that is probably a massive read if you have tens (or even hundreds) of thousands of records.

This situation may be likely to occur where searchable intelligence has been imbedded into the data content of certain fields. Usually, this is not a problem, but here are two examples where it can cause the massive SDIDS read problem:

– Using the first character of the RNID field (when it is user assigned), and then searching on RNID/L..

– Putting the area code into the phone number fields, as mentioned above, and searching on PH/804..

- **Sorting large SRLs on some fields:** Normally, sorting an SRL does not severely impair performance, even if there are a large number of hits. Sorting SRLs, however, may cause massive numbers of SDIDS records to be read in some cases. Sorting reports always uses an external sort method that never needs to read the SDIDS, and instead it always reads the SDDS.

  For search transactions that require sorting Information/Management chooses between two sort techniques and usually the result is a very fast, low-cost sort. If, however, the number of hits is large (see SORTPFX N3 parameter in the *Planning and Installation Guide and Reference*, SC34-4603 manual) and the field being used as a sort field has massive numbers of different values, then performance can be degraded by the sort.

  A common example of massive SDIDS read requirements from a sort is when large SRLs are sorted on RNID in a large database. This sort will cause *all* of the SDIDS RNID records (one for each SDDS record) to be read, until the last hit has been located. Usually this is almost *all* of the RNID records in the SDIDS. The effect is very similar to the range searches described above.

A.1.4 Is Our Performance Impacted by Massive SDIDS Reads?

Here are some items to consider that may indicate that your Information/Management users may be causing large amounts of potentially avoidable work with inefficient search criteria or sorting. If you suspect that any of the following might be true, and you are concerned about Information/Management performance, you can investigate and determine whether the recommendations in the next section will help you improve performance. Remember that inappropriate searches can come from interactive transactions, API transactions, Terminal Simulation Panels (TSPs), Report Format Tables (RFTs) and user exits:

- Do any of your Information/Management users sort SRLs on fields that have *many* thousands of different values (RNID, RNOR, etc.)?

  HINT: Check your sort fields against your glossary; a few hundred or even a few thousand different values are not a problem, but if you find your sort fields have tens or hundreds of thousand different values, then you are likely to be causing massive numbers of SDIDS reads that are hurting your performance.

- Do any of your Information/Management users use very large specific ranges in search criteria or do they include periods or asterisks that may cover many thousands of different values? The most obvious case is searching on RNID/., but other situations may also present a problem. For example, if you keep the area code in the phone number fields and you search on PH/805. to restrict a search to California then you must read one SDIDS record for *each* record that contains any California phone number in order to satisfy the search criteria.

- Do some users report long response times for searches even during periods or relatively light loads?
- Do you see CPU utilization spikes during searches?
- Do some users appear to be locked out (have much longer than normal response times) while others do searches or reports?

If the answer to any of these questions is yes, and you wish to improve Information/Management performance, then here are some things to do that will help verify whether massive SDIDS reads are really a problem in your situation:

- If you are using Version 5 or 6, you can use the QUERY function in the BLX Service Provider to count the number of GETS to the SDIDS before and after a suspect transaction. Enter the QUERY transaction from the CONSOLE before and after entering the suspect transaction and the difference in the GET count will be the number of GETS that occurred during that time. You can do this when no other Information/Management users are active to get a precise answer, but since you are looking for massive numbers of SDIDS GETS, a little interference from other users is probably transparent.
- RMF or a similar tool can help confirm the diagnosis. If you can shorten the measurement window and measure one (or several if necessary) of the suspected transactions, you should see a CPU spike that lasts throughout the transaction.
- SDSF DA can be used to measure the CPU cost for an individual transaction. You may be able to see the CPU spike at the total CPU utilization level or you can calculate the delta CPU-seconds spent on a suspect transaction.
- If you are measuring during production with many Information/Management users, you may use RMF or similar tools to see SDIDS ENQUEUE contention, since other Information/Management users will almost certainly encounter serious enqueue delays during any inappropriate search or sort activity.
- Use the procedure based on the Search Warning Parameter, as described in 7.3, “Analyzing the Application” on page 63.

A.1.5 What Can I Do?

If you determine that inappropriate searches or sorts are impairing your Information/Management performance, then there are many things you can do. Here are some of them:

- Do not sort SRLs at all, or limit the users who can. If you don′t need to sort SRLs for some users, then not sorting is an easy way to avoid the cost.
- Do not use sort fields that contain massive numbers of different values. If you must sort, then choose sort fields that do not have tens of thousands of different values. If you can sort on fields with relatively few different values, then you can even make your sorts much faster by lowering SORTPFX N3 to cause the efficient SDIDS sort to be used on even small SRLs.
- Use SORTPFX to your advantage. If you must sort on fields that may have massive numbers of different values, then use the SORTPFX parameters to decrease the probability of (or completely eliminate) inappropriate sorts.

– Increasing SORTPFX N3 will cause SRL sorting to be done from SDDS reads instead of from SDIDS reads for larger SRLs. This can reduce SDIDS contention.

– Lowering SORTPFX N2 can help you prevent inappropriate sorts by encouraging smaller SRLs, which won′t use the SDIDS for sorting. SORTPFX N2 will prevent sorting at all for SRLs with more than SORTPFX N2 ″hits″.

– Lowering SORTPFX N1 can also encourage smaller SRLs. It terminates any search that exceeds SORTPFX N1 ″hits″.

Tip

If you have many complex TSP that do searches as part of their process, be cautious with the SORTPFX command, because it could cause the TSPs to malfunction due to then hitting the limit, thus no SRL. A unique session parameter may be the answer in this situation.

HINT: You can completely eliminate SDIDS range reads for sort purposes by forcing the SDDS to be used instead. Just set SORTPFX N2 as small as your users can tolerate and set SORTPFX N3 to the same value. There will never be an SRL with more than SORTPFX N3 hits that need to be sorted, since SORTPFX N2 will have prevented the sort. Shorter SLRs will have a minimum sort cost of about 1 second per 100 hits in the SRL, so you can′t afford to make N2 and N3 too large; however, that may still be a lot better than reading 100,000 SDIDS records from the RNID range even though the smaller SDIDS records can be read as fast as 1000 per second. Another advantage of this technique is that it takes the pressure off the (usually) busier SDIDS enqueue resource and puts it on the SDDS, where you can relieve the pressure by using a five-cluster SDDS if necessary.

- Change your sort field to one that has fewer different values. For example, instead of sorting on RNID, you may be able to sort on DATE and get a search results list that is not exactly in RNID order, but close enough to satisfy your need.

- If you can′t find a suitable field to sort on, you may be able to *make* a suitable field and then sort on that new field. You can use the WORDFIX feature to create a new field in all of your records that has a copy of the ″search intelligence″ part of the original field. You do not even have to remove the original data. Then you can search on the new field very quickly, since there will only be a few SDIDS records related to it.

  For example, if you created a new AREACOD field using the three digits copied from the PH/ field, then your searches (and sorts) on this new field would give the same result as the inappropriate big-range searches required from the original field.

  Another example of using WORDFIX to avoid massive SDIDS reads involves creating a new field in each of your SDDS records that concatenates the data from the record′s DATE and TIME fields. Then you can sort on that field instead of using an RNID/. search. If you use system assigned RNIDs or if your own RNID assignment technique assigns RNIDs sequentially, then a search on the new combined DATETIME field will result in the same hits in RNID order, but you will not have to read as many SDIDS records to do the sort. It probably won′t reduce the number of SDIDS records in the sort range much; you will still have one SDIDS record for each DATETIME where any new records were filed. That may still be a lot, maybe almost as many as the comparable RNID range for the records.

But there is one thing that can save a lot using this technique compared to sorting on RNID. Since the records in a hit list are most likely to be relatively recent records they will have a higher DATETIME value since Information/Management stops looking up the sort order in the SDIDS when all the hits have been found. You can take advantage of this fact by ″inverting″ the DATETIME field by subtracting it from 9999999 with a user exit before filing it. Then when you search on problem records and sort by DATETIME the sort is cut short as soon as the last hit has been found in the look-up table. Since the look-up table (not the sort) is inverted, then the sort will finish much faster when all the records are relatively recent. It won′t help if even one hit is very old, but if all hits are relatively recent, then many thousands of SDIDS reads can be avoided.

- Discourage the use of freeform search arguments:
  - Use QUICK SEARCH instead of freeform arguments.
  - Use TSP controlled searches instead of freeform arguments.

  Make sure you check your TSPs, SRCs, and QUICK SEARCH to make sure they do not cause massive search reads.

- Eliminate unnecessary big-range searches:
  - Check your TSPs, APIs, RFTs and user exits for inappropriate range searches.
  - Educate your users who must do freeform searches. Explain how they can avoid massive SDIDS reads by not doing big-range searches.

- In Version 5, there is one other parameter that can help. SRCHENQ determines how often a user who is resolving search criteria will release the enqueue that it must get during the search reads. This parameter defaults to 0, which means: ″Don′t release until search processing is finished reading each range of SDIDS records″. During the enqueue most other users who need the SDIDS to proceed with their transactions will be blocked. If you set SRCHENQ to a positive integer, then the enqueue will be released after many SDIDS records in each search range have been read. That will prevent the search user from completely blocking other Information/Management transactions, but in busy situations, it will lengthen the response time of the inappropriate search. Here are some guidelines on setting SRCHENQ to an appropriate value:

  In Version 6, the function provided by SRCHENQ has been enhanced and called SRCHLIMIT. It is a three-value parameter. The first value works like SRCHENQ, the second value allows you to send a warning to any user that exceeds a specified number of SDIDS range reads, and the third value allows you to terminate the search or sort, when the number of SDIDS range reads has been exceeded. You can use these new features of the SRCHLIMIT parameter to help you identify and avoid transactions that are using excessive CPU resources.

  - *Never* set SRCHENQ in V5 or the first parameter of SRCHLIMIT in V6 to a number less than 10, because continuing the read operation after each SRCHENQ operation requires the position in the SDIDS to be re-established, thereby increasing the total number of reads. Thus setting SRCHENQ to 10 can add one extra read for each 10 originally required.

- Setting SRCHENQ to 100 only adds 1% to the total I/O count, but would provide frequent opportunities for other users to get access to the SDIDS. This will tremendously improve many response times, especially for transactions that only require a few enqueue opportunities.
- Other competing transactions that may require many SDIDS enqueue opportunities (for example, filing a new record typically requires 50 to 150 SDIDS enqueue opportunities) may be improved with an even lower SRCHENQ but you will be increasing the effect of the ″reread-the-last-record″ situation. If you wish to favor users filing new records (in their competition with range search users) then you might consider setting SRCHENQ lower than 100, maybe even as low as 20. At that level, you will increase the number of I/Os required of the search user by 5%, but you will give the create user many more opportunities to get the SDIDS enqueue.
- Larger SRCHENQ values (200-500) can be effective for many transactions but will not help large transactions such as filing new records very much.

• V5 customers should also apply the APAR fix for OY62656; it causes SRCHENQ to correctly release the enqueue for the SDIDS during SORT reads. It also adds a new parameter (COGENQ) that allows you to prioritize create, update, copy, or purge activity by user. COGENQ is also in V6. COGENQ allows the cognize activity to hold the SDIDS enqueue while several (COGENQ) SDIDS records are cognized. Setting COGENQ to the default (10) for a user means that when a user files a record the transaction will only cognize 10 fields before releasing the enqueue, whereas a V5 user without the fix would hold the enqueue until all fields have been cognized, thus holding up other users. This raises that user′s priority relative to other users of the SDIDS, such as range searchers or range sorters or even relative to other ″file″ users with lower COGENQ settings. COGENQ saves a little CPU, but more importantly, it means users filing records get *faster* response times if they are competing with SDIDS read abusers.

A.1.6 Summary

Prior to Version 5, abusive search criteria or expensive sort situations were bad, but tolerable. Their effects (primarily CPU utilization and SDIDS enqueue utilization) are worse in Versions 5 and 6. Response times for the abusive transaction can be much longer, and response times for competing users can also be increased.

To get maximum benefits from the significant function and performance improvements that Version 5 and 6 offer, you should avoid the abusive search criteria and sort situations described here. Where very large range reading of the SDIDS cannot be avoided, it is important to take maximum advantage of available tuning parameters.

Version 4 performance can also benefit (though not as much) from avoiding abusive search criteria or sort situations.

Appendix B. Performance Recommendations Checklist

The checklist that follows is intended as a quick reference of the key areas that this redbook has been discussing:

- Use 18 or 34-byte SDIDS keys to eliminate spanned SDIDS records.
- If you are not using Multisystem Database Access, check your GRS parameters in case any enqueues are being treated as global when they shouldn′t be. If possible, avoid using multiple-cluster SDDSs; if not, minimize the number of sort-related reads on the SDDS.
- Carefully work out your VSAM buffer values to gain maximum benefit.
- Enable Advanced SDIDS Compression.
- If possible, use cache and DASD FastWrite for your SDDS and SDIDS data sets.
- Use VSAM buffers and PNLBCNT to minimize the I/O to your read panel data sets.
- Eliminate unnecessary cognizing by uncognizing s-words.
- Consider the physical placement of your VSAM files to minimize contention.
- Utilize LPA for Information/Management code.
- Avoid STEPLIB and JOBLIB where possible.
- Minimize the concatenations for STEPLIB and JOBLIB.
- Minimize the concatenations LLA.
- Consider raising the dispatching priority of your Information/Management users.
- Avoid large searches.
- Use COGENQ, SORTPFX and SRCHLIMIT to tune the interleaving of enqueues.
- Avoid Multisystem Database Access on the same MVS machine.
- If are using Multisystem Database Access, then review the values for the GRS parameters RESMIL and ACCELSYS (or the equivalent parameters if a different product is being used).

If you are running a version previous to Version 6.3, you may want to take into consideration the following:

- If your SDDS has less than 261 000 records, it should be defined as NONSPANNED, and have a CISIZE sufficient to contain the largest record.
- If your SDDS has more than 261 000 records, then you will have to define your SDIDS as SPANNED. Choose a large CISIZE: 32768 will usually be appropriate to minimize the number of CIs required to hold the few very long SDIDS records, one being the master bit list.

Appendix C. TSP With S-word

The TSP included in this appendix shows an example of how to use s-word for searching. You should never use p-word in TSP or Reports to look for a specific record type (for example, STAC/.). Using s-word instead of p-word reduces the number or accesses to SDIDS data sets.

```
+--------------------------------------------------------------------------+
| PURYEARL                                                             PMF |
|                                                                          |
| With this TSP all service records are searched                           |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
+--------------------------------------------------------------------------+
```

TERMINAL SIMULATOR PANEL COMMON CONTROL FLOW DATA

LINE NUM	FUNCTION NAME	LABEL NAME	LITERAL/TEST DATA	PANEL NAME	FUNCTION EXIT	FIELD NAME	WORD OCCUR	APPLY NOT	GET VAR
1	MOVEVAR								
2	ADDDATA		TABLE MC629291, SE						YES
3	ADDDATA								NO
4	ADDDATA		,TABLE BLG1TSRL,END						NO
5	PROCESS	ERROR							
6	RETURN								
7	LABEL	ERROR							
8	RETURN								

TERMINAL SIMULATOR PANEL COMMON CONTROL DATA

LINE NUM	FUNCTION NAME	S-WORD INDEX	STRUCTURED WORD	PREFIX INDEX	PREFIX WORD	PREFIX VALIDATION PATTERN	VERIFY NAME	VERIFY TYPE	USER DATA
1	MOVEVAR								
2	ADDDATA	0000							

PMF PANEL CONTENT REPORT

DATE 22.12.95 PANELS: PURYEARL

TIME 12:33 DATA SET LABEL: RPANEL4

TERMINAL SIMULATOR PANEL COMMON CONTROL DATA

LINE NUM	FUNCTION NAME	S-WORD INDEX	STRUCTURED WORD	PREFIX INDEX	PREFIX WORD	PREFIX VALIDATION PATTERN	VERIFY NAME	VERIFY TYPE	USER DATA
3	ADDDATA	0B19	X//QCSESC						
4	ADDDATA	0000							
5	PROCESS								
6	RETURN								
7	LABEL								
8	RETURN								

TERMINAL SIMULATOR PANEL UNIQUE CONTROL DATA

LINE NUM	FUNCTION NAME	NEW S-WORD INDEX	NEW S-WORD	NEW PREFIX INDEX	NEW PREFIX	NEW PREFIX VALIDATION PATTERN	NEW DATA
1	MOVEVAR						
2	ADDDATA						
3	ADDDATA						
4	ADDDATA						
5	PROCESS						
6	RETURN						
7	LABEL						
8	RETURN						

```
TERMINAL SIMULATOR PANEL UNIQUE FLAG DATA

LINE  FUNCTION     PRINT  PRINT  PRINT  CURR  LAST  RETAIN    TRACE    TRACE    FIND STRING  STRING  SAVE GEN  DATA
NUM     NAME       MSGS   SCREN  TSCA   REC   REC   REC ID     ON      LINK      ANYWHERE     DATA   MESSAGE   TYPE
---  ----------    ---    ---    ---    ---   ---    ---       ---      ---         ---        ---      ---     ----

  1  MOVEVAR                                                                                                   CHAR
  2  ADDDATA
  3  ADDDATA
  4  ADDDATA
  5  PROCESS                                                                                         NO
  6  RETURN
  7  LABEL
  8  RETURN
TERMINAL SIMULATOR PANEL COMMON CONTROL LINE

STARTING  COMM   DIALOG    PANEL    ---------- LEVEL --------- ----- LAST ALTERED ----
 PANEL    PANEL  SECTION  MODIFIED    USER    REL PTF FMID/APAR   DATE    TIME    USER
--------  ---      --       ---     -------- --  --  --------  -------- ----- --------

           NO      00       YES               00  00            30.11.94 15:21 N2P1BN2
```

Appendix D. Special Notices

This redbook is intended to help systems engineers and systems programmers who need to tune performance and do capacity planning in Information/Management. The information in this publication is not intended as the specification of any programming interfaces that are provided by Information/Management. See the PUBLICATIONS section of the *Programming Announcement for TME 10 Information/Management* for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM′s product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM′s intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter 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 the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM (″vendor″) products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer′s ability to evaluate and integrate them into the customer′s operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

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

Advanced Function Printing	AFP
AIX	APPN
CBIPO	CICS
CICS/ESA	CT
CUA	DATABASE 2
DB2	DFSMS/MVS
DFSORT	Enterprise Systems Connection Architecture
ESCON	GDDM
IBM	Information Warehouse
MVS/DFP	MVS/ESA
MVS/SP	NetView
OfficeVision/VM	OPC
Operating System/2	OS/2
OS/390	PS/2
RACF	RETAIN
RS/6000	SAA
SP	SystemView
System/390	VisualAge
VTAM	WebExplorer

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

Java and HotJava are trademarks of Sun Microsystems, Incorporated.

Microsoft, Windows, Windows NT, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

Pentium, MMX, ProShare, LANDesk, and ActionMedia are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

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

C + +	American Telephone and Telegraph Company, Incorporated
Display Manager	Digital Research, Incorporated
HAL	Qantel
Lotus Notes	Lotus Development Corporation
Netscape	Netscape Communications Corporation
Sun, Solaris	Sun Microsystems, Incorporated
Windows NT	Microsoft Corporation
80386	Intel Corporation

Other trademarks are trademarks of their respective companies.

Appendix E. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

E.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see "How to Get ITSO Redbooks" on page 99.

- *Information/Management V6.3. Network Computing and Workstation Interfaces*, SG24-4574
- *Information/Management V6.3 Installation as an Enterprise-Wide Server Web Connector, E-Mail and Remote API Support*, SG24-4573
- *Boosting Your Information/Management Version 6 Release 1 Performance*, GG24-4324
- *Jumping out of the Box: New Interfaces and Features of Information/Management Version 6 Release 1*, GG24-4325
- *Installing Information/Management Version 6 Release 1 for Client/Server*, GG24-4323

E.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
Application Development Redbooks Collection	SBOF-7290	SK2T-8037
Personal Systems Redbooks Collection	SBOF-7250	SK2T-8042

E.3 Other Publications

These publications are also relevant as further information sources:

- *Planning and Installation Guide and Reference*, SC34-4603
- *Terminal Simulator Guide and Reference*, SC34-4609
- *Client Installation and User's Guide*, SC34-4593
- *CICS/ESA Systems Definition Guide*, SC33-0664
- *CICS/ESA Resource Definition (Online)*, SC33-0666

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at http://www.redbooks.ibm.com.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **PUBORDER** — to order hardcopies in United States
- **GOPHER link to the Internet** - type GOPHER.WTSCPOK.ITSO.IBM.COM
- **Tools disks**

  To get LIST3820s of redbooks, type one of the following commands:

  ```
  TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
  TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
  ```

  To get BookManager BOOKs of redbooks, type the following command:

  ```
  TOOLCAT REDBOOKS
  ```

  To get lists of redbooks, type one of the following commands:

  ```
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LISTSERV PACKAGE
  ```

  To register for information on workshops, residencies, and redbooks, type the following command:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1996
  ```

  For a list of product area specialists in the ITSO: type the following command:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
  ```

- **Redbooks Web Site on the World Wide Web**

  http://w3.itso.ibm.com/redbooks

- **IBM Direct Publications Catalog on the World Wide Web**

  http://www.elink.ibmlink.ibm.com/pbl/pbl

  IBM employees may obtain LIST3820s of redbooks from this page.

- **REDBOOKS category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL
- **Internet Listserver**

  With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibmlink.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

Redpieces

For information so current it is still in the process of being written, look at ″Redpieces″ on the Redbooks Web Site (http://www.redbooks.ibm.com/redpieces.htm). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** — send orders to:

	IBMMAIL	Internet
In United States:	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada:	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America:	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+45) 4810-1170 - Swedish

- **Mail Orders** — send orders to:

IBM Publications	IBM Publications	IBM Direct Services
Publications Customer Support	144-4th Avenue, S.W.	Sortemosevej 21
P.O. Box 29570	Calgary, Alberta T2P 3N5	DK-3450 Allerød
Raleigh, NC 27626-0570	Canada	Denmark
USA		

- **Fax** — send orders to:

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States)** or **(+1)001-408-256-5422 (Outside USA)** — ask for:

  Index # 4421 Abstracts of new redbooks
  Index # 4422 IBM redbooks
  Index # 4420 Redbooks for last six months

- **Direct Services** - send note to softwareshop@vnet.ibm.com
- **On the World Wide Web**

Redbooks Web Site	http://www.redbooks.ibm.com
IBM Direct Publications Catalog	http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Internet Listserver**

  With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibmlink.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank).

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (http://www.redbooks.ibm.com/redpieces.htm). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

- Invoice to customer number

- Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

List of Abbreviations

ABEND	Abnormal end
AMS	Access Method Services
AFP	Advanced Function Printing
APA	all points addressable
APF	Authorized program facility
API	Application program interface
APPC	Advanced Program-to-Program Communication
CAS	central address space
CGI	Common Gateway Interface
CLIST	Command List
CM/2	Communication Manager/2
CICS	Customer Information Control System
CISIZE	Control interval size
CTC	Channel to channel
DFSORT	Data Facility Sort
DGA	Database Gateway Application
DLC	Data link control
DLL	Data link library
ESCON	Enterprise Systems Connection Architecture
e-mail	electronic Mail
FRP	Forms Processing Routines
GUI	Graphical User Interface
GRS	Global Resource Serialization
HICA	High-level application communication area
HLAPI	High-level application program interface
HTML	Hypertext Markup Language
ID	Identifier
IETF	Internet Engineering Task Force
IRC	Immediate Reponse Chain
IBM	International Business Machines Corporation
JCL	Job control language
ISPF	Interactive System Productivity Facility
ITF	Installation Tailoring Facility
ITSO	International Technical Support Organization
LAN	Local Area Network
LLAPI	Low-level application program interface
LSR	Local shared resources
LRECL	Local record length
LU	Local unit
MRES	Multiclient remote environment server
MSDA	Multisystem Database Access
MVS	Multiple Virtual Storage
NCCF	Network Communications Control Facility
NCP	Network Control Program
NDF	Node Definition File
OS/2	Operating System/2
OPC/A	Operating Planning and Control/Advanced
OPC/ESA	Operating Planning and Control/Enterprice Systems Architecture
OS/390	Operating System/390
PCF	Panel Control Facility
PDBs	parameter data blocks
PDF	Program Development Facility
PDS	Partitioned data set
PIAT	Program interface argument table
PICA	Program interface communication area
PIDT	Program interface data table
PIHT	Program interface history table
PIMB	Program interface message block
PIPT	Program interface pattern table
PIRT	Program interface result table
PMF	Panel Modification Facility
PROFS	Professional Office System
PWS	Programmable workstations
PU	Physical Unit

P-word	prefix word
RAB	Record anchor block
RACF	Resource Access Control Facility
RDO	Resource definition (online)
RES	Remote environment server
RFC	Request for Comments
RFT	Report format table
RGV	REXX Global Variable
RMF	Resource Measurement Facility
RMODE	Residency mode
RNID	Record number identifier
RPANLDS	Read panels data set
SDDS	Structure description data set
SDE	Structure description entry
SDIDS	Structure description index data set
SDLC	Synchronous Data Link Control
SDLDS	Structure description log data set
SDR	Structure description record
SDSF	Spool Display and Search Facility
SLR	Service Level Reporter
SMP	System Modification Program
SMTP	Simple Mail Transfer Protocol
SNA	System Network Architecture
SRCs	Store response chains
SSN	System name
S-word	Structured word
TP	Transaction program
TSCA	Terminal simulator communication area
TSO	Time sharing option
TSP	Terminal simulator panel
TSX	Terminal simulator EXEC
URL	Uniform Resource Locator
URI	Uniform Resource Indicator
VSAM	Virtual storage access method
WPANLDS	Write panel data set

Index

T

U

V

W

ITSO Redbook Evaluation

TME 10 Information/Management Version 1.1 Performance and Capacity Considerations
SG24-4571-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at http://www.redbooks.com
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@vnet.ibm.com

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction ____________

Please answer the following questions:

Was this redbook published in time for your needs? Yes____ No____

If no, please explain:

____

____

____

____

What other redbooks would you like to see published?

____

____

____

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

____

____

____

____

____

Printed in U.S.A.

SG24-4571-00