

RMAN Performance Tuning Using Buffer Memory Parameters

Introduction

The purpose of RMAN performance tuning is to identify the bottlenecks for a given backup or restore job and use RMAN commands, initialization parameters, or adjustments to physical media to improve overall performance.

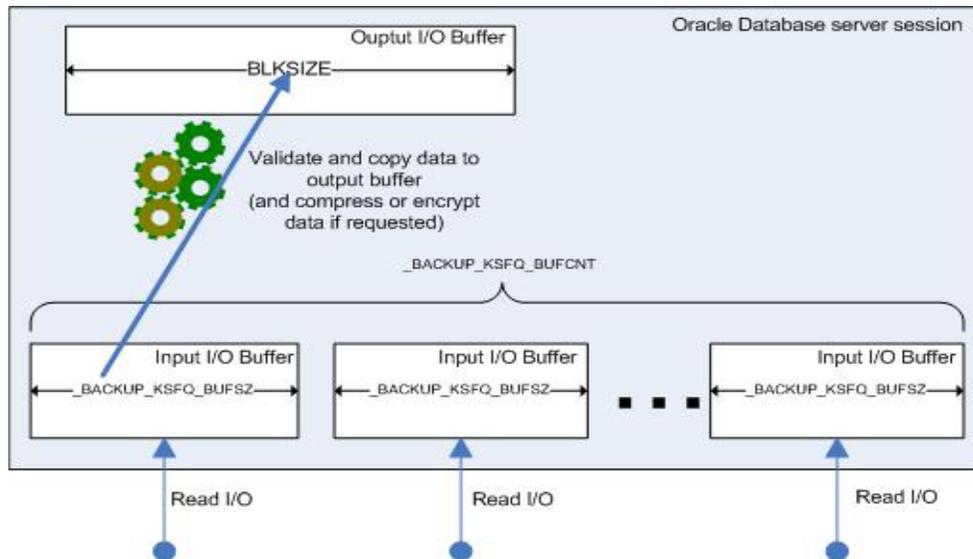
As database sizes continue to grow unabated, with tens to hundreds of terabytes prevalent in customer's environments, server and storage systems have also dramatically increased in resource, network, and I/O capacity to support required database backup and restore windows. Against this backdrop, RMAN performance tuning becomes more important than ever.

This document details RMAN performance tuning guidelines as they relate to buffer memory usage during backup and restore operations. These guidelines are targeted to systems where the available I/O exceeds 100 MB/s, while memory, CPU, network consumption, and tape drive throughput do not constrain overall RMAN performance. This is especially true in very high performance configurations, such as Exadata. In these environments, the default RMAN buffer sizes are usually inadequate and should be increased.

It is assumed that the reader has a basic understanding of RMAN performance tuning concepts. See this [doc](#) for an overview of RMAN tuning.

RMAN Buffer Usage

During backup, datafile blocks are read into a set of input buffers, where they are validated/compressed/encrypted and copied to a set of output buffers. The output buffers are then written to backup pieces on either disk or tape (`DEVICE TYPE DISK` or `SBT`). This process is depicted in the below diagram.



Conversely, on restore, blocks are read from the backup piece into a set of input buffers, where they are validated/uncompressed/decrypted and copied to a set of output buffers. The output buffers are then written back to data files on disk.

By default, buffers are allocated from the PGA, unless DISK or TAPE I/O slaves are used, in which case buffers are allocated from the SGA or large pool (if `LARGE_POOL_SIZE` is set).

The number and size of the buffers used for each RMAN job can be found in the views `V$BACKUP_ASYNC_IO` and `V$BACKUP_SYNC_IO`.

The following guidelines for setting RMAN memory usage are organized by Oracle release.

Tuning guideline:

- When increasing buffer size using any of the following underscore parameters, it is recommended to increase both input and output buffer sizes simultaneously for a given operation.

Oracle Database 11g Release 11.1.0.7 and prior releases

When creating backup sets:

Input buffers are allocated based on the level of *multiplexing* (i.e. number of files read per channel), according to the below chart:

Multiplexing = Min (MAXOPENFILES, FILESPERSET)

Defaults: MAXOPENFILES=8, FILESPERSET=64

Number of Files Read Per Channel	Buffer Size
Multiplexing ≤ 4	Each buffer = 1MB, total buffer size for channel is up to 16MB. Number of buffers per file will depend on the number of files.
4 > Multiplexing ≤ 8	Each buffer = 512KB, total buffer size for channel is up to 16MB. Numbers of buffers per file will depend on number of files.
Multiplexing > 8	Each file will have 4 512KB buffers.

For example, if `MAXOPENFILES=1`, each channel uses 16 buffers, each of size 1MB, for a total of 16 MB allocated per channel.

Tuning guideline:

- The number and size of input buffers can be changed by using the undocumented parameters `_backup_ksfq_bufcnt` and `_backup_ksfq_bufsz`.
- When using these parameters to increase backup I/O throughput, it is recommended that `_backup_ksfq_bufsz` be first set to the stripe size of the files being read and to test the effect that setting has on performance.
- If further tuning is needed, then set `_backup_ksfq_bufcnt` to the number of disks which the files are striped over. In the case of massively striped systems, such as ASM, this parameter should be used judiciously to achieve an optimal balance between performance and memory usage.
 - For example, if there are 500 disks in your ASM diskgroup, setting `_backup_ksfq_bufcnt` to 500 would use an excessive amount of memory for each RMAN channel. A more reasonable setting, such as 32 or 64, provides a better balance between performance and memory usage.

When creating or restoring from image copies:

The size and number of input buffers default to 1 MB/buffer and 4 buffers.

Tuning guideline:

- Input buffer sizes for image copies can be increased using `_db_file_direct_io_count`.

When creating *backup sets* or *image copies*:

Output buffers are allocated based on device type:

Device Type	Number of Buffers Per Channel	Buffer Size	Total of Buffers Allocated
DISK	4	1MB	4MB
SBT	4	256KB	1MB

Tuning guideline:

- For DISK backup, the output buffer size can be increased using `_db_file_direct_io_count`.
- Note that `_db_file_direct_io_count` controls other database I/O operations, not just RMAN, and therefore its adjustment should be carefully tested.
- For SBT backup, the output buffer size can be increased using `BLKSIZE` channel parameter.
- Note that `BLKSIZE` does not apply to Oracle Secure Backup. In this case, buffer sizes are automatically calculated for optimal transfer speeds.

When restoring from *backup sets*:

The default size and number of output buffers is 128 KB/buffer and 4 buffers.

Tuning guideline:

- The number of output buffers can be increased using `_backup_ksfq_bufcnt`.
- The output buffer size can be increased using `_backup_ksfq_bufsz`.
- Note that the default buffer sizes for restoring files are smaller than what are used when backing up those same files. These buffers should be increased to match the corresponding backup if restore times are significantly slower than the corresponding backups.

When restoring from *image copies*:

The default size and number of output buffers is 1 MB/buffer and 4 buffers.

Tuning guideline:

- Output buffer size can be increased using `_db_file_direct_io_count`.

Oracle Database 11g Release 11.1.0.7 + Patch 8369105 (also in Exadata V1 Bundle Patch 2)

Input and output buffers are allocated and tuned largely the same as in the preceding section, with the exception of the following changes:

- When restoring from *backup sets*, the default output buffer size is increased from 128 KB to 1 MB.
- `_db_file_direct_io_count` is not used to adjust any RMAN buffer sizes. This means that RMAN buffer sizes can be adjusted without affecting any other database I/O.
- `_backup_ksfq_bufcnt` can be used to adjust the number of input and output buffers for all operations using DISK channels.

Oracle Database 11g Release 2

All RMAN I/O operations with ASM files select an optimal buffer size and count based on the allocation unit (AU) size and disk count of the ASM disk group.

Additional parameters are introduced to give more fine-grained control over RMAN buffer allocations. Although these parameters are supported for all RMAN operations, they are primarily intended for use in non-ASM systems, since RMAN operations in ASM will automatically select an optimal buffer configuration.

- `_backup_seq_bufsz` / `_backup_seq_bufcnt`
 - These parameters set the buffer size and count for both creating and restoring backup pieces using SBT channels with third-party media managers.
 - Note that these parameters are not needed when using Oracle Secure Backup, as the optimal buffer size is automatically selected.

- `_backup_disk_bufsz/_backup_disk_bufcnt`
 - These parameters set the buffer size and count for both creating and restoring backup pieces using DISK channels.
- `_backup_file_bufsz/_backup_file_bufcnt`
 - These parameters set the input buffer size and count for all backup operations, and the output buffer size and count for all restore operations.
- Note that `_backup_ksfg_bufsz` and `_backup_ksfg_bufcnt` are still supported in Oracle Database 11g Release 2 for compatibility reasons, but it is recommended that you tune the RMAN buffer sizes instead using the above parameters.