



Smarten

Advanced Data Discovery

Powered by ElegantJ BI

The Workings of Managed Memory Computing

Business Intelligence & Advanced Data Discovery

Document Information	
Document ID	Smarten-The-Workings-of-Managed-Memory-Computing
Document Version	3.0
Product Version	5.0 and above
Date	15-December-2018
Recipient	NA
Author	EMTPL

© Copyright Elegant MicroWeb Technologies Pvt. Ltd. 2018. All Rights Reserved.

Statement of Confidentiality, Disclaimer and Copyright

This document contains information that is proprietary and confidential to EMTPL, which shall not be disclosed, transmitted, or duplicated, used in whole or in part for any purpose other than its intended purpose. Any use or disclosure in whole or in part of this information without the express written permission of EMTPL is prohibited.

Any other company and product names mentioned are used for identification purpose only, may be trademarks of their respective owners and are duly acknowledged.

Disclaimer

This document is intended to support administrators, technology managers or developers using and implementing Smarten. The business needs of each organization will vary and this document is expected to provide guidelines and not rules for making any decisions related to Smarten. The overall performance of Smarten depends on many factors, including but not limited to hardware configuration and network throughput.

Contents

1	About this Document	4
2	Managed Memory Computing – Overview.....	4
3	Understanding Managed Memory Computing	4
4	Managed Memory Computing - Working	5
4.1	Terminologies	5
4.2	General Guideline for Memory Allocation	6
4.3	Lifecycle of the Data Cube or dataset in Managed Memory Computing	6
5	The Decisions for Managed Memory Computing	7
6	Product and Support Information	7

1 About this Document

This document will enable the users of Smarten to understand the concept of Managed Memory Computing. This document also describes the workings of Managed Memory Computing.

It is recommended that this document be read after reading Smarten Hardware Sizing document.

Note:

Throughout this document, Smarten Advanced Data Discovery is abbreviated as Smarten.

2 Managed Memory Computing – Overview

Smarten introduces a unique concept of Managed Memory Computing which delivers fast response time and high performance by effectively utilizing hybrid technology.

There are two technical approaches available in market today. The first is a cube or dataset based approach where data cubes and datasets remain on disk and processing is dependent on disk I/O operations. The second is the in-memory approach where data is brought into the memory for processing.

Managed Memory computing allows technology managers to configure Smarten for the use of dual technology by combining both approaches, i.e. the administrator has the option to keep some data cubes or datasets in-memory while retaining others on the disk. Once the data cube or dataset is in Managed Memory, user analytics and other requests will be computed in-memory and will not involve any significant disk I/O. The data cube or dataset will stay in-memory and will be refreshed automatically when cube or dataset is refreshed.

3 Understanding Managed Memory Computing

Managed Memory Computing is different from conventional in-memory processing as well as from cube or dataset based computing.

Conventional in-memory processing does not provide a logical business layer and does not aggregate data before it is brought in-memory. Cube or dataset based processing provides the power of aggregated data and logical business layer. The traditional power of aggregated data cubes or datasets, which provides a logical business layer and allows dimensions and measures to be set, remains available in Managed Memory Computing.

Managed Memory Computing, like in-memory processing, retains the data in the memory until explicitly removed or replaced by other data. In Managed Memory Computing, data cubes loaded in-memory do not include the flat files of source data or result sets. The data cubes or datasets brought in-memory are comprised of pre-processed data, optimized for high performance.

Managed Memory Computing makes disk I/O optional and not a rule. Traditional cube based analysis tools retain large cubes on disks, resulting in increased disk I/O which in turn results in deteriorating response time for users.

Managed Memory Computing does not take an either-or approach but instead allows the administrator of the technology to choose which data should be kept in-memory and which needs to be accessed from the disk. It is not necessary that all the data cubes or datasets be selected for Managed Memory Computing, thus providing a very powerful option to administrators for allocating hardware resources based on priority of response times.

4 Managed Memory Computing - Working

Overall memory requirements are discussed at length in the Hardware Sizing document of Smarten and a short summary is provided below for quick reference.

The memory on the server in a typical scenario will be used for the components below:

- Operating System
- Application Server
- Smarten objects which are being accessed by users
- Smarten cubes or datasets which have been kept in-memory (based on administrator selection)
- Usage of memory by any other applications or services which are not for operating Smarten

Memory is allocated to Smarten based on memory configuration in an application server such as JBOSS. It is highly recommended that the server is devoted to Smarten and resources are not shared with any other application when considering the use of Managed Memory Computing.

4.1 Terminologies

Terminology used in Managed Memory Computing

- **Allocated Managed Memory:** Memory allocated to Managed Memory Computing at the current moment.
- **Used Managed Memory:** Portion of the memory being used out of the total memory allocated for Managed Memory Computing at the current moment.
- **Free Memory:** The difference between Allocated and Used Memory at the current moment.
- **Allocation for Managed Memory Computing:** The amount of memory to be allocated for Managed Memory Computing. This is also a setting in the administrative interface. Once saved this becomes the Currently Allocated Managed Memory.

Determining the amount of Managed Memory is based on the recommended Managed Memory and Total Available Memory as well as the overall managed memory settings planned and operational for other data cubes or datasets.

Ideally, the total Allocated Managed Memory should be greater than the sum of the size of each cube or dataset selected for Managed Memory. This eliminates the possibility of forced ejection of a cube or dataset and ensures the consistent high performance of all selected cubes or datasets.

- **Total Available Memory:** The maximum memory allocated in the application server. In JBOSS this is the xmax setting.
- **Recommended Allocation for Managed Memory:** The recommended allocation is 60% of the maximum available memory. This recommendation is based on general assumptions, and may not apply to specific configurations or needs.
- **Current Managed Memory Requirements:** Memory required for Managed Memory Computing based on the data cubes and datasets which have been selected for Managed Memory.

As indicated earlier, the Hardware Sizing document and JBOSS Optimization Guide is to be referred to for determining the amount of physical RAM required in the server.

4.2 General Guideline for Memory Allocation

- Application Server -maximum memory setting
 - As a general practice, adequate memory is first allocated for the smooth running of the Operating System. The balance of the memory can be allocated to the Application Server. This will be the maximum setting in the Application Server.
- Application Server- minimum memory setting
 - The minimum memory allocation in the Application Server will be higher when you are using Managed Memory Computing. As a standard, the minimum should be 70% of the maximum memory allocation described above.
 - Having a very low value as a minimum memory allocation is not prudent when using Managed Memory Computing as it will cause frequent resizing of the allocated memory each time new data needs to be processed in the memory.

By default, none of the data cubes or datasets are enabled for Managed Memory Computing. The administrator must explicitly select cubes or datasets for Managed Memory Computing.

It is important to note that the memory allocated for Managed Memory Computing is independent of the memory used by user objects.

Managed Memory Computing for a data cube or dataset and allocation of managed memory in Smarten is configured through the administrator interface.

4.3 Lifecycle of the Data Cube or dataset in Managed Memory Computing

The memory allocated to Managed Memory Computing is dedicated solely to the activity of managing data cubes or datasets in-memory. Sessions in-memory for user objects such as dashboards and analysis do not have any impact on this memory, and memory for these user objects is allocated by the Application Server as per available memory and settings. If the memory allocated for Managed Memory Computing is not used by the data cubes or datasets, the Application Server will use it for other processes and tasks, including for user objects.

Initiation of the Data Cube or dataset into the Memory

- The data is brought into memory when a user requests an object which needs data from a particular data cube or dataset which is selected for Managed Memory Computing and that cube or dataset is not already loaded into the memory.
- All the data of a cube or dataset is brought into Managed Memory.

In-memory Operations

- When other objects request data from the cube or dataset already in memory, the user will receive a very quick response as the data cube or dataset is already in memory.

Memory Management and Forced Ejections

- The memory is freed based on the LRU (Least Recently Used) scheme..
 - Least Recently Used (LRU)

When memory needs to be freed, the data elements which last used timestamp is smaller, will be removed from the memory. If the memory allocated for Managed Memory Computing has been fully utilized and new data needs to be loaded into the memory, the data of the least frequently used cube or dataset will be unloaded from the memory.

- If the data from a cube or dataset is being used and hence cannot be unloaded from memory to accommodate new data, the requested data will be provided by reading from the disk.
- If the user is accessing data from a cube or dataset in memory, the refresh process of a rebuilt cube or dataset will be put on hold until all users have stopped accessing the data cube.
- When the cube or dataset is being refreshed into the memory, the user will be put on hold until the relevant data cube or data set is loaded into the memory.

Any changes made by the administrator, such as a change in security settings at the column level, will come into effect only when the user makes a fresh request for the data for which the settings have been changed. Security settings will apply irrespective of whether the cube or dataset is in-memory or not.

5 The Decisions for Managed Memory Computing

The decision to keep a data cube or dataset in-memory or to allow access from the disk is dependent primarily on the business needs and the allocation of hardware resources, especially RAM.

When business needs dictate that all analysis should be available at the best speeds, this does not mean that an infinite amount of memory is available for in-memory processing. Managed Memory Computing will require the system administrator to allocate priority for data cubes or datasets based on the following:

- Keep data cubes or datasets which are accessed once or twice a day on conventional disk I/O mode.
- Select data cubes or datasets which are accessed regularly and are critical for quick analysis for Managed Memory Computing.
- Choose to put data cubes or datasets in-memory for presentations or meetings where there is a requirement for granular data review for strategic or tactical decision making. These cubes or datasets can be removed from Managed Memory Processing after the meeting and discussion is complete. It is imperative that the Allocated Memory for Managed Memory Computing be adjusted during these temporary changes by deselecting the appropriate data cubes or datasets out of Managed Memory Computing.

6 Product and Support Information

Find more information about Smarten and its features at www.smarten.com

Support: support@smarten.com

Sales: sales@smarten.com

Feedback & Suggestions: support@smarten.com

Support & Knowledgebase Portal: support.smartent.com