



ORACLE® ESSBASE

Release 11.1.2.2

New Features

ORACLE®
ENTERPRISE PERFORMANCE
MANAGEMENT SYSTEM

This document describes the new features in Oracle Essbase Release 11.1.2.2. For detailed information on these features, see the *Oracle Essbase Database Administrator's Guide*, *Oracle Essbase Technical Reference*, and *Oracle Essbase API Reference*.

For information about the new features in Oracle Essbase Administration Services, see *Oracle Essbase Administration Services New Features*.

CONTENTS IN BRIEF

Block Storage Performance Enhancements	2
Optimization of Distinct Count Calculations on Aggregate Storage Databases	3
Distinct Way to Reference Shared Members in MDX	4
Changing Essbase Server and Application Log Levels in Administration Services	4
MDX Property Expressions	4
Throttle Active Transactions	4
Essbase API Changes	5
Configuration Setting (essbase.cfg) Changes	5
SET CALCPARALLEL Calculation Command Enhancement	6
New MDX Functions	6
MaxL Grammar Changes	6

Block Storage Performance Enhancements

Subtopics

- [Block Storage Parallel Calculation](#)
- [Block Storage Parallel Restructuring](#)
- [Block Storage Parallel Data Export](#)
- [Block Storage Parallel Data Load](#)
- [Optimizing Block Storage Parallel Data Load](#)
- [Cache Sizes Larger than 4 GB](#)

Block Storage Parallel Calculation

Calculation tasks are generated along the last n sparse dimensions of an outline. These sparse dimensions used to identify tasks are called task dimensions. The number of task dimensions, n , is either selected dynamically by Essbase, or you can override the number by specifying a value for `CALCTASKDIMS` in the `essbase.cfg` file.

Block Storage Parallel Restructuring

By default, block storage restructuring is performed sequentially. Blocks are renumbered and reshaped from first to last, a time-intensive process. Parallel restructuring can reduce restructuring time by dividing block restructuring work across multiple concurrent threads to use available processor cores. Because calculation is performed separately from restructuring, each block can be restructured independently of other blocks.

The number of restructuring threads to use is defined in `essbase.cfg` using the `RESTRUCTURETHREADS` configuration setting.

Block Storage Parallel Data Export

Using the **export data** MaxL statement to export data from a block storage database in parallel, you specify a comma-separated list of export files, up to a maximum of 1024 file names. The number of file names determines the number of export threads. The number of available block-address ranges limits the number of export threads that Essbase actually uses. Essbase divides the number of actual data blocks by the specified number of file names (export threads). If there are fewer actual data blocks than the specified number of export threads, the number of export threads that are created is based on the number of actual data blocks. This approach results in a more even distribution of data blocks between export threads.

Note: In specifying the number of export files, it is important to consider the number of available CPU cores and I/O bandwidth on the computer on which Essbase Server runs. Specifying too large a number can result in poor performance.

The EXPORTTHREADS configuration setting has been enhanced to support block storage parallel data export. (EXPORTTHREADS is used only when executing the PAREXPORT ESSCMD command without specifying the -threads option.)

Block Storage Parallel Data Load

Parallel data load refers to the concurrent loading of multiple data files into an Essbase database. When working with large data sets (for example, a set of ten 2 GB files), loading the data sources concurrently enables you to fully utilize the CPU resources and I/O channels of modern servers with multiple processors and high-performance storage subsystems.

Parallel data load uses multiple parallel pipelines on the server side, and multiple threads on the client-side, to load multiple data files concurrently, thus enabling data loads to be truly optimized to the capabilities of modern servers.

Optimizing Block Storage Parallel Data Load

Parallel data load includes a “throttle” to limit the number of client threads or server pipelines that are used. A throttle control is important because if the data-file specification matches several hundred files, you should avoid spawning that many threads or pipelines. To control the number of threads or pipelines spawned by a data load request, set the limit using the **max_threads** grammar in the **import data** MaxL statement.

Cache Sizes Larger than 4 GB

The 64-bit Essbase Server has been expanded to accommodate larger index, data, and data file cache sizes without setting a scaling factor. The MEMSCALINGFACTOR configuration setting is deprecated. For full database functionality, you must upgrade to the latest client software.

For more information, see the Essbase Readme for this release.

Optimization of Distinct Count Calculations on Aggregate Storage Databases

On aggregate storage databases only, the NonEmptyCount MDX function is optimized so that the calculation of the distinct count for all cells can be performed by scanning the database only once. Without this optimization, the database is scanned as many times as the number of cells corresponding to the distinct count. The NonEmptyCount optimization is triggered when an outline member formula has the following syntax:

```
NONEMPTYCOUNT(set, measure, exclude_missing)
```

The **exclude_missing** parameter supports the NonEmptyCount optimization on aggregate databases by improving the performance of a query that queries metrics that perform a distinct count calculation.

Distinct Way to Reference Shared Members in MDX

In an outline that is enabled for duplicate member names, all member names, including shared members, are distinguishable by using a qualified member name syntax. However, in a unique member name outline (where duplicate members are *not* enabled), a member and its shared member are both identified by the same simple member name.

In previous releases, any reference in an MDX query to such a shared member was assumed to be a reference to the original member, not the shared member. With this release, you can get a unique name for shared members by using the MEMBER_UNIQUE_NAME property in an MDX query.

For example, in the Sample Basic database, the member [100-20] is an original member under parent [100], and has a shared member associated with it under parent [Diet]. The shared member [100-20] can be referred to explicitly, using the unique name [Diet].[100-20], as shown in the following query:

```
SELECT
  {Sales}
ON COLUMNS,
  {[Diet].[100-20]}} PROPERTIES MEMBER_UNIQUE_NAME
ON ROWS
FROM Sample.Basic;
```

For rules on the syntax and escape characters required to uniquely reference shared members in a unique member name outline, see the topic entitled “MDX Syntax for Specifying Duplicate Member Names and Aliases” in the *Oracle Essbase Technical Reference*.

Changing Essbase Server and Application Log Levels in Administration Services

You can change Essbase log levels at the server and application level in Administration Services Console. See “Changing Log Levels” in the *Oracle Essbase Administration Services Online Help*.

MDX Property Expressions

In addition to querying for intrinsic and custom properties of a particular member, you can also query for properties of related members, using the PROPERTY_EXPR function within the DIMENSION PROPERTIES section of the query axis.

See the MDX Property Expressions topic and the CurrentAxisMember topic in the *Oracle Essbase Technical Reference*

Throttle Active Transactions

Each allotted server thread for an application may create child threads for tasks such as parallel calculation, parallel data load or export, and parallel restructuring. However, if the total number

of running threads is too high, threads may lose efficiency while contending for server resources. To control the issue of thread contention, you can use the `MAXTOTALACTIVETRANSACTIONS` and `MAXACTIVEUPDATETRANSACTIONS` `essbase.cfg` settings.

Essbase API Changes

Subtopics

- [Changed C Main API Structures](#)
- [Compiler Support](#)

Changed C Main API Structures

The data structures `ESS_DBINFO_T` and `ESS_DBSTATE_T` have expanded fields to accommodate 64-bit long typedefs. This is related to [“Cache Sizes Larger than 4 GB”](#) on page 3.

Compiler Support

The supported compilers list for building Essbase API programs has been updated. See the “Supported Compilers” topic in the *Oracle Essbase API Reference*.

Configuration Setting (essbase.cfg) Changes

Subtopics

- [New Configuration Settings](#)
- [Deprecated Configuration Setting](#)
- [CALCPARALLEL Configuration Setting Enhancement](#)

New Configuration Settings

- `MAXACTIVEUPDATETRANSACTIONS`. See [“Throttle Active Transactions”](#) on page 4.
- `MAXTOTALACTIVETRANSACTIONS`. See [“Throttle Active Transactions”](#) on page 4.
- `NUMBLOCKSTOEXTEND`: Determines the number of bytes by which data files in block storage databases are extended to accommodate block updates that require additional disk space.
- `RESTRUCTURETHREADS`: Specifies whether parallel restructuring is enabled for a database, and the number of threads to use.
- `XOLAPMAXINCLAUSESIZE`: Sets the maximum number of `INLIST` items allowed in SQL clauses (for XOLAP databases).

Deprecated Configuration Setting

The MEMSCALINGFACTOR configuration setting is deprecated. The 64-bit Essbase Server has been expanded to accommodate larger cache sizes without setting a scaling factor.

CALCPARALLEL Configuration Setting Enhancement

The maximum number of threads that can be made available for parallel calculation has been increased:

- Block storage databases:
 - Running on 32-bit platforms: 64
 - Running on 64-bit platforms: 128
- Aggregate storage databases (running on 32-bit or 64-bit platforms): 128

SET CALCPARALLEL Calculation Command Enhancement

The maximum number of threads that can be made available for parallel calculation has been increased:

- 32-bit platforms: 64
- 64-bit platforms: 128

New MDX Functions

- PROPERTY_EXPR function (MDX Property Expressions), for querying properties of ancestor members.
- CurrentAxisMember, for use with the PROPERTY_EXPR function.

MaxL Grammar Changes

- The **query database** MaxL statement for aggregate storage databases has a new **list aggregate_storage slice_info** clause, to display information about data slices and views.
- The **deploy** MaxL statement for deploying cubes from Oracle Essbase Studio has the following new grammar:
 - If an Oracle Essbase model is enabled for XOLAP, you can, when redeploying the cube, use the **outline_in_background** grammar.
 - When the model being deployed contains single or multiple relational data sources, you can deploy in streaming mode using the **use streaming build** grammar if you want Oracle

Essbase Studio to query the external data source directly (rather than using an ODBC connection).

- The **import data** MaxL statement has a new **using max_threads** clause, for specifying maximum number of threads to use for a parallel data load.
- The **alter system** and **alter application** MaxL statements have new **set ... message_level** clauses, to set the message logging level at the agent and application scopes.
- The **display system** and **display application** MaxL statements have new **message_level** clauses, to display the current message level settings.

COPYRIGHT NOTICE

Essbase New Features, 11.1.2.2

Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Authors: EPM Information Development Team

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS:

Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.