

# Using Precompute for Better Performance

## SAS OpRisk Monitor 3.4.1 with hot fix 4

The precompute configuration and pre-processing is a new capability for improved response in systems that have a high volume of dimensions, mappings, and mapping constraints. This capability mainly affects the ORA chooser component, yet might also improve degraded processing performance related to assessment processing in the SAS OpRisk Monitor web application. The use of this new configuration option is suggested if you meet the following criteria:

1. Database installation is Oracle. Oracle 9 and Oracle 10 / 10G are supported.
2. many thousands of dimensions loaded into your business structure.
3. many thousands of mappings relating the components of your business structure
4. multiple constraints configured in dimensionality.xml
5. “constrained by mappings” check box selected in the ORA chooser.

The following SAS notes detail possibilities where pre-compute could be a strategy to use:

**32823: SAS® OpRisk monitor web application performs slowly**

<http://support.sas.com/kb/32/823.html>

**32845: Out of memory errors when loading large volumes of data - possible data corruption**

<http://support.sas.com/kb/32/845.html>

**32859: Adding a user to notify when creating an assessment produces no results**

<http://support.sas.com/kb/32/859.html>

**32912: How to improve assessment performance using parameter monitor.riskAssesments.answerSheets.ignoreRoles**

<http://support.sas.com/kb/32/912.html>

**32860: No response when attempting to view assessors for a questionnaire that has many points in its area**

<http://support.sas.com/kb/32/860.html>

After applying SAS OpRisk Monitor 3.4.1 hot fix 4, the ORA chooser supports multiple methods for computing which dimension nodes to present when the “constrain by mappings” option is enabled.

The original method for this computation, prior to hot fix 4, analyzed the relevant constraints and computed the available dimension nodes dynamically. This method is now referred to as the “no-cache” option and is the default configuration.

The new method uses a batch process to analyze the relevant constraints and to pre-compute the valid combinations. This new method is referred to as the “pre-compute” option.

Use the new configuration option, `monitor.ora.chooser.strategy`, to specify the method (no-cache or pre-compute). The new configuration options and a new command line utility added in hotfix 4 are discussed below.

To apply this configuration, the general steps are:

- a. Decide if you want to apply index changes to your Oracle database
- b. Modify the `configdata.properties` file and set the appropriate configuration options
- c. Run `precompute.cmd` (windows) or `precompute.sh` (unix) which are located in the SAS OpRisk Monitor Administration `dbscripts` directory.
- d. Re-deploy the monitor web application after making the changes (including the database index changes).

## I. New Configuration Options

### monitor.ora.chooser.strategy

Supported value: pre-compute or no-cache (the default).

Use no-cache option if you have a small number of mappings and number of constraints and small number of users and security definitions. This option retains the same ORA chooser behavior as the previous OpRisk releases.

Use pre-compute if your database is in Oracle 9 or 10/10G, and you have a large number of mappings (thousands), a large number of dimensions (thousands), multiple constraints configured in `dimensionality.xml`, and you have the “constrained by mappings” check box selected in the ORA chooser.

The following is a comparison and suggestion on when to use which options.

	Pre-computation	No-Cache
Pros	<ul style="list-style-type: none"> <li>▪ Apply to the large number of mappings</li> </ul>	<ul style="list-style-type: none"> <li>▪ Allow the changes of the mappings and the user's security in real time.</li> <li>▪ Provide a solution to be consistent with the previous release</li> </ul>
Cons	<ul style="list-style-type: none"> <li>▪ Manual steps to reflect the update of the constraints / mappings</li> <li>▪ Does not recognize changes to the mappings and the user's security in real time.</li> </ul>	<ul style="list-style-type: none"> <li>▪ Can not apply to large number of mappings</li> <li>▪ When the mappings/security information is updated, users may experience a delay when using ORA chooser the first</li> </ul>

		time.
Conditions to use (the numbers are configurable)	<ul style="list-style-type: none"> <li>▪ Number of mappings for any feature &gt; 10,000</li> <li>▪ Number of security constraints for any user &gt; 200</li> </ul>	<ul style="list-style-type: none"> <li>▪ Number of mappings for any feature &lt; 10,000</li> <li>▪ Number of security constraints for any user &lt; 200</li> </ul>

The new pre-compute method may return slightly different results than the original no-cache method. In most cases, the pre-compute method eliminates “dead ends” that prevent a user from making a selection for all the required dimensions or is eliminating invalid combinations when multiple nodes are selected for a single dimension. The details of these differences are listed below.

- A. The pre-compute method analyzes the full set of constraints at one time. The no-cache strategy will only analyze the relationships that include the current requested dimension. In certain situations, the no-cache will present a node for selection because it is valid for the considered mappings but the user will hit a dead end and will not be able to find a valid node for a subsequent dimension.
- B. The pre-compute method requires that all mapping constraints defined for a feature **must** have overlapping dimensions. If there are no overlapping dimensions, then the pre-compute method decides that there is no intersection between the constraints. For example if the Management Org - Legal Org Map and the Risk Cause Map were used as the constraints on a feature, there would be no overlapping dimensions. The no-cache method decides that all combinations are valid.
- C. If the selected dimension is not part of the dimensions in both mapping and security constraints, and there is no intersection between security and mappings, the listed message “no data to display” is shown. If there is at least one intersection, all data is displayed. The no-cache method will display all available nodes.
- D. To determine which dimensions participate in each mapping, the pre-compute method looks at the actual data in the database rather than the list of dimensions specified in the dimensionality.xml file. This allows the pre-compute strategy to perform the dimensions analysis using current data -- even if the information in dimensionality.xml is out-of-date. The no-cache method relies on the information in dimensionality.xml. If the information in dimensionality.xml matches the database, the two strategies will determine that the same dimensions participate in the mappings.
- E. In the case, where multiple nodes are chosen for a single dimension, the two methods analyze differently. For example, if two management organizations are selected, the no-cache method displays any node that was mapped to either of the two management organizations. The pre-compute method displays only the nodes that are mapped to both of the two management

organizations. Because the current system builds an operational risk area out of all combinations of the chosen nodes, the no-cache method can build an area that includes points that are not valid combinations.

F. Parent constraints are ignored in pre-compute mode.

### **monitor.ora.chooser.preselect.strategy**

Supported value: all or any

This config option specifies how the pre-selected nodes will impact the selection of other dimensional nodes. "All" is for the "AND" relationship and "any" is for the "OR" relationship. Only two combinations are supported as listed below.

	pre-compute	no-cache
All	Yes	No
Any	No	Yes

### **monitor.ora.chooser.recomputeOnLogin**

Supported value: true or false

Setting this value to true will clear the user's cached ORA chooser constraint information on login. The default value is false because enabling this feature can cause fragmentation of the user\_survivor\_node table. If user data is only updated periodically via batch, then user information should be cleared as part of that batch process. If user information changes frequently, then this setting should be set to true, but user\_survivor\_node table should be cleared and truncated on a regular basis.

## **II. New Command Line Utility**

A new command line utility called "precompute" has been added to the admin tools. Precompute uses the information in file dimensionality.xml to compute all the valid dimension node combinations for all application features (all ORA chooser configurations) and store those combinations in the database. Before running the precompute command, be sure the copy of your dimensionality.xml in the administration tools classes directory is in synch with the one in the WEB-INF/classes directory of the deployed web application on the web-tier.

Note: the pre-compute batch process analyzes only the mapping data on features that have multiple constraints (2 or more). You should re-run precompute whenever mappings used in a dimensionality feature constraint are changed.

It is strongly recommended that the database is backed up before running precompute.

Each time you follow the steps to implement the pre-compute method, you must delete and recreate the database index REL\_INTR\_AK1, and then analyze table RELATION\_INTERSECTION:

To delete the database index, issue the SQL command:

```
drop index REL_INTR_AK1;
```

Note: It is possible the index does not exist yet. As a result you may get a message "Error: ORA-01418: specified index does not exist". This is not a problem and you may proceed.

To recreate the database index, issue the SQL command::

```
CREATE INDEX REL_INTR_AK1 ON  
RELATION_INTERSECTION(INTERSECTION_TXT,MANAGEMENT_ORG_RK,PROC  
ESS_RK, RISK_CAT_RK, PRODUCT_RK,BL_RK, RESOURCE_DIM_RK,  
AUX_RMG_DIM_1_RK, AUX_ORG_DIM_1_RK, AUX_OP_DIM_1_RK,  
LEGAL_ORG_RK, COST_CENTER_RK, GEOGRAPHY_RK,  
PROJECT_RK,OBJECTIVE_RK, CONTROL_RK, CAUSE_RK, NFI_TYPE_RK,  
AUX_RISK_DIM_1_RK, BASEL_BL_RK, BASEL_RISK_CAT_RK, STD_PROC_RK,  
FSI_RK, AUX_RPT_DIM_1_RK);
```

To analyze, issue the SQL command:

```
Analyze table RELATION_INTERSECTION compute statistics;
```

For optimal performance (assuming the OpRisk schema name is *lossdata*):

- 1) Delete and recreate Oracle statistics using the following SQL command, each time after running precompute:

```
exec dbms_stats.delete_schema_stats(ownname => 'lossdata');  
exec dbms_stats.gather_schema_stats(ownname => 'lossdata', method_opt =>  
'for all columns', cascade => true, degree => dbms_stats.DEFAULT_DEGREE ,  
estimate_percent => 100);
```

- 2) Periodically refresh statistics on table USER\_SURVIVOR\_NODE using the following SQL command:

```
exec dbms_stats.gather_table_stats(ownname => 'lossdata', tablename =>  
'user_survivor_node', method_opt => 'for all columns', cascade => true, degree  
=> dbms_stats.DEFAULT_DEGREE , estimate_percent => 100);
```

- 3) Deactivate the configured issuing of Oracle ANALYZE commands by SAS OpRisk Monitor data loaders. In file <ADMIN\_COMPONENT\_INSTALL\_DIR>\Tools\3.2\classes\configdata.properties where the SAS OpRisk Monitor Administration component is installed, remove com.sas.oprisk.framework.server.dataload.listeners.OrderOfMagnitudeAnalyzeListener from the dataload.listeners property.

NOTE: Up to date database statistics are optimal. Database administrators should schedule regular statistics collection and always recreate or restore statistics when restoring the database/schema from backup. At no time should monitor be running without statistics having first been collected. It is expected that statistics might only be re-gathered on a weekly basis for some customers and that would be deemed sufficient.

### Additional Oracle Configuration

1. Disable the recycle bin – the following command should be run once by the database administrator.

```
ALTER SYSTEM SET RECYCLEBIN=OFF;
```

2. Keep the recycle bin disabled - the following parameter should be set in the Oracle initialization file.

```
RECYCLEBIN=OFF;
```

3. Set the following Oracle initialization parameter to reduce the time spent parsing the many SQL statements issued in the assessment process.

```
CURSOR_SHARING=SIMILAR;
```

4. Provide enough CPU and memory for concurrent assessments and heavy ORA chooser activity. The more cpu's and memory that are available the better. We recommend that the database administrator **start** with a 4GB SGA, 1GB PGA, and 8GB of system memory. They should these values if they see heavy cpu or memory utilization on their hardware. For example:

4GB to 8GB SGA

16-32 GB system memory

Also the PGA should be increased if warnings are given via enterprise manager.

Note: Assessment creation can be resource intensive when there is a large number of points (several hundred) in the assessment's operational risk area. Heavy ORA chooser use happening concurrently with assessment creation might be burdensome on the system. In this situation, it may be useful to suggest a time window for assessment creation.