

CA Unified Infrastructure Management

Probe Guide for Data Engine
data_engine v8.0 series



Copyright Notice

This online help system (the "System") is for your informational purposes only and is subject to change or withdrawal by CA at any time.

This System may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This System is confidential and proprietary information of CA and protected by the copyright laws of the United States and international treaties. This System may not be disclosed by you or used for any purpose other than as may be permitted in a separate agreement between you and CA governing your use of the CA software to which the System relates (the "CA Software"). Such agreement is not modified in any way by the terms of this notice.

Notwithstanding the foregoing, if you are a licensed user of the CA Software you may make one copy of the System for internal use by you and your employees, provided that all CA copyright notices and legends are affixed to the reproduced copy.

The right to make a copy of the System is limited to the period during which the license for the CA Software remains in full force and effect. Should the license terminate for any reason, it shall be your responsibility to certify in writing to CA that all copies and partial copies of the System have been destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS SYSTEM "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IN NO EVENT WILL CA BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS SYSTEM, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF CA IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The manufacturer of this System is CA.

Provided with "Restricted Rights." Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

Copyright © 2014 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Legal information on third-party and public domain software used in this product is documented in the *Third-Party Licenses and Terms of Use* (http://docs.nimsoft.com/prodhelp/en_US/Library/Legal.html).

Contact CA

Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information that you need for your Home Office, Small Business, and Enterprise CA Technologies products. At <http://ca.com/support>, you can access the following resources:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Providing Feedback about Product Documentation

Send comments or questions about CA Technologies product documentation to nimsoft.techpubs@ca.com.

To provide feedback about general CA Technologies product documentation, complete our short customer survey which is available on the support website at <http://ca.com/docs>.

Contents

Chapter 1: Overview	9
How the data_engine Probe Collects and Maintains QoS Data	9
RN_QOS_DATA Table Columns	10
RN_table Indexes	10
Data_engine Start Up	10
Parallel Mode	11
Serial Mode	12
Chapter 2: Configuration Tasks	13
Prerequisites for SQL Authentication on Microsoft SQL Server	13
Preparing to Partition Your Oracle Database	14
Upgrade Prerequisites for Deployments with an Oracle Database	16
Change the Database Connection Properties	17
Configure the Data Retention Settings.....	17
Override the Data Retention Settings on Individual QoS Objects	18
Set up Index Maintenance (MS SQL Server and Oracle)	18
Set up Partitioning for Raw Sample Data (MS SQL Server and Oracle)	19
Schedule Database Maintenance	19
Chapter 3: Best Practices	21
Chapter 4: Configuration Details	23
data_engine	23
Database Configuration.....	24
MySQL	25
Microsoft SQL Server	26
Oracle.....	28
Quality of Service	29
Scheduler.....	30
Chapter 5: Troubleshooting	31
Viewing the Log File	31
Corrupted QoS Definition Values	32
Out of date information in the Quality of Service Type Status Table	34
Table Statistics for Oracle	34

Table Statistics for MySQL.....	35
Table Statistics for Microsoft SQL Server	36
Check Partitioning Jobs	38

Documentation Changes

This table describes the version history for this document.

Version	Date	What's New?
8.0	September 2014	Updated to include partitioning and index maintenance for Oracle.
7.9	June 2014	Added information for daily tables and Advanced Configuration items.
7.9	March 2014	Added information for new GUI items.
7.9	March 2013	Initial release of Admin Console probe GUI. (Previous versions of this probe are only configured using Infrastructure Manager).

Related Documentation

Documentation for other versions of the data_engine probe

The [Release Notes](#) for the data_engine probe

Monitor Metrics Reference Information for CA Unified Infrastructure Management Probes

(http://docs.nimsoft.com/prodhelp/en_US/Probes/ProbeReference/index.htm)

Chapter 1: Overview

The `data_engine` probe manages and maintains data that is collected by Quality of Service (QoS) enabled probes. The `data_engine` creates all tables and stored procedures necessary to manage the collected data.

Data that is produced by the QoS probes is stored in the UIM database, in tables named raw data tables (table prefix is `RN_`). Raw data is kept for a user-defined period, then summarized and aggregated into Hourly data (`HN_tables`). Hourly data is then summarized and aggregated into Daily data (`DN_tables`).

How the `data_engine` Probe Collects and Maintains QoS Data

During data collection, `data_engine` performs the following actions:

1. The `data_engine` probe receives a QoS definition from another probe.
2. The `data_engine` probe queries the `S_QOS_DEFINITION` table to determine if the QoS type in the message already exists in the table (for example, `QOS_CPU_USAGE`).
3. If the QoS type does not exist, a new entry is added to the `S_QOS_DEFINITION` table. New `RN`, `HN`, `BN` and `DN` tables are created in the UIM database to store monitoring data from the probe.
4. When the first `qos_message` from the probe arrives, the `data_engine` probe adds QoS object data into the `S_QOS_DATA` table. The `S_QOS_DATA` table contains the object data for each unique combination of `qos`, `source`, and `target` attributes.
5. The `data_engine` probe inserts the raw data from the probes into the appropriate `RN` tables.
6. During the scheduled maintenance runs, the `data_engine` probe summarizes and aggregates raw data from `RN` tables into hourly data that is stored in `HN` tables.
7. Older `RN` data is purged based on a user-defined period.
8. The `data_engine` probe then summarizes and aggregates hourly data from `HN` tables into daily data that is stored in `DN` tables.
9. Older `HN` data is purged based on a user-defined period.
10. Last sample value coming from the probes for each `qos` object is populated in the `S_QOS_SNAPSHOT` table. This data is used to provide fast QoS data access for UMP portlets.

RN_QOS_DATA Table Columns

RN_QoS_Data_tables hold raw QoS data. QoS data is written once and never updated.

Column Name	Description
TableID	unique identifier; key for looking up time series data
Sampletime	time the sample was taken
Samplevalue	QoS value
Samplestdev	standard deviation of the sample
Samplerate	Rate of sampling
Tz_offset	time zone offset

RN_table Indexes

The default indexes in RN_tables are optimized for writing data:

Index	Description
Idx0	table_id, sampletime
Idx1	table_id, sampletime, samplevalue, samplerate, tz_offset

The RN_QoS_DATA_tables do not have primary keys, as both tableID and sampletime can be duplicated.

Data_engine Start Up

When the data_engine probe starts, it loads both S_QOS_DATA and S_QOS_DEFINITION into memory and establishes a bulk connection to the database.

The data_engine probe has two modes:

- [Parallel Mode](#) (see page 11)
- [Serial Mode](#) (see page 12)

Parallel Mode

Note: We recommend running the data_engine probe in parallel mode to improve data_engine probe performance.

If the thread_count_insert parameter is set to a number greater than one in the Raw Configure menu, the data_engine probe uses parallel mode.

In parallel mode, the data_engine uses multiple threads to do work in parallel. In this mode, the data_engine:

1. Continuously reads data from the hub.
2. Stores the data in a shared list and begins reading more messages.
3. Iterates over all lists using another thread to see if any of the lists need to be flushed to the database.

Note: The actual writing of data does not happen now.

4. All objects are marked as 'ready to commit' and a reference is placed onto another list.
5. Concurrently, a thread continuously runs to validate, sort, and place messages into a list that is written to the database.
6. Writes any messages that have been marked as 'ready to commit data' to the database using a thread pool of worker threads.

This process allows the data_engine probe to write many rows to the database in parallel. Also, the data_engine is not as vulnerable to performance issues when inserting into one table is slow.

For more information about setting parallel mode, see [Best Practices](#) (see page 21).

Serial Mode

Serial mode is the original mode for the data_engine probe. If the thread_count_insert parameter is set to zero in the Raw Configure menu, the data_engine probe defaults to serial mode.

In this mode, the data_engine:

1. Reads messages from the hub for a given time period (default around 1 second).
 - a. From 1 through 20 messages are read at a time, depending on how many are in the queue or if the hub queue size has changed (default 20).
 - b. The read messages are then validated and sorted into lists that can be quickly inserted into the database.
2. Stops reading new messages and iterates over all the lists, checking to see if any are full. By default the list is flushed if it contains more than 5000 messages or if it has not been flushed in the last 5 seconds.
3. Goes back to reading messages from hub. If one bulk object takes too long to insert, then the writing of all data to the database is delayed.

Chapter 2: Configuration Tasks

If you are using SQL Authentication on Microsoft SQL Server, review the [Prerequisites for SQL Authentication on Microsoft SQL Server](#). (see page 13)

Review [Preparing to Partition Your Oracle Database](#) (see page 14) to ensure that you have enough free disk space to create a partitioned interim table that is used during the partitioning process.

Before upgrading UIM Server 8.0 in deployments with an Oracle database, you must grant the appropriate user permissions. See [Upgrade Prerequisites for Deployments with an Oracle Database](#) (see page 16) for details.

In the Admin Console data_engine configuration GUI, you can:

- Change the Database Connection Properties.
- Configure the Data Retention Settings.
- Override the Data Retention Settings on Individual QoS Objects.
- Schedule Database Maintenance.
- Set up Partitioning for Raw Sample Data (MS SQL Server)
- Set up Index Maintenance (MS SQL Server)

To open the data_engine configuration GUI:

1. In the Admin Console navigation tree, click the down arrow next to the hub, then the robot the data_engine probe resides on.
2. Click the down arrow next to the data_engine probe, select **Configure**.

Prerequisites for SQL Authentication on Microsoft SQL Server

If you are not using the System Administrator (sa) login and require SQL Authentication on Microsoft SQL Server, your user account must have the following permissions:

- The db_owner database role for the UIM database.
- Read and update permissions on the master and tmpdb system databases.
- The serveradmin database role to create and execute stored procedures properly.

Preparing to Partition Your Oracle Database

Although partitioning your Oracle database provides enhanced performance, be aware that depending on the size and number of data tables in your database, the partitioning process can take a few minutes, several hours, or possibly several days and can be resource intensive. In addition to the processing time required to complete partitioning, the partitioning process requires the availability of enough free disk space to create a partitioned interim table. The partitioning process copies the contents from the original data tables into the interim data tables, swaps the original data tables and interim data tables, and finally deletes the original data tables. This process is performed while the tables and indexes are online. If there is not enough free disk space for the interim tables, the partitioning process will not be able to finish. The process is similar to the one described in this [Oracle article](#).

This section provides a list of tasks we recommend you perform before selecting the **Partition Data Tables** option.

1. The partitioning process requires the default tablespace (for the interim table) to be 1.6 times larger than the largest data table to be partitioned. In general, RN tables are the largest data tables in the database. To find the largest RN table, generate and run the following script:

```
select
  segment_name      table_name,
  sum(bytes)/(1024*1024) table_size_meg
from
  user_extents
where
  (segment_type='TABLE' or segment_type='TABLE PARTITION')
and
  segment_name like 'RN_QOS_DATA_%'
group by segment_name
order by table_size_meg desc;
```

2. Initial partitioning of the tables online also consumes disk space in the SYSTEM tablespace. The disk space required prior to partitioning is approximately the same as that required by the largest unpartitioned RN table. Determine the utilization of the system tables by generating and running the following script:

```
column dummy noprint
column pct_used format 999.9 heading "%|Used"
column name format a19 heading "Tablespace Name"
column Kbytes format 999,999,999 heading "KBytes"
column used format 999,999,999 heading "Used"
column free format 999,999,999 heading "Free"
column largest format 999,999,999 heading "Largest"
column max_size format 999,999,999 heading "MaxPoss|Kbytes"
```

```

column pct_max_used format 999.9 heading "%|Max|Used"
break on report
compute sum of kbytes on report
compute sum of free on report
compute sum of used on report

select (select decode(extent_management,'LOCAL','*'') ||
        decode(segment_space_management,'AUTO','a ','m ')
        from dba_tablespaces where tablespace_name = b.tablespace_name) ||
nvl(b.tablespace_name,
    nvl(a.tablespace_name,'UNKOWN')) name,
kbytes_alloc mbytes,
kbytes_alloc-nvl(kbytes_free,0) used,
nvl(kbytes_free,0) free,
((kbytes_alloc-nvl(kbytes_free,0))/
    kbytes_alloc)*100 pct_used,
nvl(largest,0) largest,
nvl(kbytes_max,kbytes_alloc) Max_Size,
decode( kbytes_max, 0, 0, (kbytes_alloc/kbytes_max)*100) pct_max_used
from ( select sum(bytes)/1024/1024 Kbytes_free,
        max(bytes)/1024/1024 largest,
        tablespace_name
    from sys.dba_free_space
    group by tablespace_name ) a,
( select sum(bytes)/1024/1024 Kbytes_alloc,
    sum(maxbytes)/1024/1024 Kbytes_max,
    tablespace_name
    from sys.dba_data_files
    group by tablespace_name
    union all
    select sum(bytes)/1024/1024 Kbytes_alloc,
        sum(maxbytes)/1024/1024 Kbytes_max,
        tablespace_name
    from sys.dba_temp_files
    group by tablespace_name )b
where a.tablespace_name (+) = b.tablespace_name
order by PCT_USED
/

```

- Initial partitioning of the tables also consumes storage in the TEMP tablespace. The free disk space required prior to partitioning is approximately the same as that required by the largest unpartitioned RN table. Determine the utilization of the TEMP tablespace by generating and running the following script:

```

SELECT tablespace_name, SUM(bytes_used)/1024/1024 M_Used,
SUM(bytes_free)/1024/1024 M_Free
FROM V$temp_space_header
GROUP BY tablespace_name;

```

Upgrade Prerequisites for Deployments with an Oracle Database

Before upgrading UIM Server 8.0 in deployments with an Oracle database, you must grant the appropriate user permissions. Grant the permissions using a tool such as Oracle SQL Developer. You must be logged in as SYSDBA. The command to grant user permissions is:

```
grant execute on dbms_redefinition to <NMS USER>;
```

If the upgrade is performed before you grant the appropriate user permissions, UIM Server errors will occur during scheduled maintenance for data_engine. The following error message appears in the data_engine log file:

```
SPN_DE_DATAMAIN is invalid
```

To correct this situation, use the following procedure to grant the appropriate permissions and recompile the stored procedures.

Follow these steps:

1. Log in to the database server as SYSDBA and execute:
grant execute on dbms_redefinition to <NMS USER>;
2. Use a tool such as Oracle SQL Developer to recompile the following stored procedures:

```
SPN_DE_DATAMAIN
```

```
SPN_DE_DATAMAINDELETEOLDDATA
```

```
SPN_DE_UNPARTITIONTABLE
```

```
SPN_DE_PARTITIONTABLE
```

```
SPN_DE_UPDATEQOSDEFMETRIC
```


Change the Database Connection Properties

Important! The database connection properties should only be changed in limited circumstances such as recovery operations. Changing the Database Vendor can cause connection issues. If you are changing database vendors, CA recommends reinstalling CA Unified Infrastructure Management (UIM).

Follow these steps:

1. In the data_engine probe configuration menu, click the **Database Configuration** folder.
2. Click the **Connection Information** drop-down list, select the Database Vendor for your database.

The **Connection Information** field displays the connection options for the selected vendor.

3. Enter the connection settings. Settings are different for each database vendor. See [Database Configuration](#) (see page 24) for more information about the fields that are required for each vendor.
4. Click the **Test Connection** button.

The **Test ADO Connection String** window appears. If the connection is good, the Connected and Ping values are set to yes.

5. Click **Save**.

The configuration changes are saved and the probe is restarted.

Configure the Data Retention Settings

You can change the data retention settings to meet your auditing or security requirements.

Follow these steps:

1. In the data_engine probe configuration menu, click the **data_engine** heading.
2. Change the desired retention settings in the General Section. See [data_engine](#) (see page 23) for more information about each field.
3. Click **Save**.

The configuration changes are saved and the probe is restarted.

Override the Data Retention Settings on Individual QoS Objects

You can override the data retention settings for individual QoS items.

Follow these steps:

1. In the data_engine probe configuration menu, click the **Quality of Service** folder.
2. In the Quality of Service Table, click the row of the QoS metric you want to modify.
3. Change the desired retention settings. See [Quality of Service](#) (see page 29) for more information about each field.
4. Click **Save**.

The configuration changes are saved and the probe is restarted.

Set up Index Maintenance (MS SQL Server and Oracle)

Important! (MS SQL Server) It is not possible to rebuild the index for single partitions prior to SQLServer 2014. You can only reorganize individual partitions. Performing automatic indexing for large tables from the data_engine is discouraged, as indexing might not complete in a reasonable amount of time.

You can set up Index Maintenance to improve the speed of data retrieval operations.

Follow these steps:

1. In the data_engine probe configuration menu, click the Database Configuration folder.
2. Select the **Index Maintenance** check box.
3. Change the desired Index Maintenance options. See [Microsoft SQL Server](#) (see page 26) or [Oracle](#) (see page 28) for more information about the individual fields for each database vendor.

Important! (Oracle) If you have partitioned your Oracle data tables, there is no need to select the Index Maintenance option.

4. Click **Save**.

Index Maintenance will be performed during the next maintenance period.

Set up Partitioning for Raw Sample Data (MS SQL Server and Oracle)

You can set up partitioning to improve performance when accessing the raw sample data tables.

Follow these steps:

1. In the data_engine probe configuration menu, click the Database Configuration folder.
2. Select the **Partition Data Tables** check box.
3. Click **Save**.

Partitioning will be performed during the next maintenance period. The time required to execute the partitioning is dependent on both the amount of data and the performance of the disk subsystem. Partitioning can take up to several days on especially large installations.

Schedule Database Maintenance

You can schedule automatic database maintenance to optimize system performance.

Follow these steps:

1. In the data_engine probe configuration menu, click the **Scheduler** folder.
2. Enter a maintenance start date. You can set the start date for a future date or can start the maintenance schedule immediately.
3. Enter a maintenance end date. The end date can either be a calendar date or a set number of occurrences.
4. Select a Recurrence pattern.

The additional time options appear for your selected duration pattern.

5. Enter the additional time options that are required for your selected duration. For example, if you select a daily duration pattern of days=6, hours=0, and minutes=0, maintenance occurs every 6th day at midnight.
6. Click **Apply**.

The configuration changes are saved. If the start now option was selected, the new schedule for maintenance begins.

Chapter 3: Best Practices

The data_engine performs most tasks with little or no interaction from the administrator. However there are some configurations that can improve performance.

data_engine Probe Location

- In order to reduce the network traffic, run the data_engine on a Hub as close to the database server as possible.

Thread Count

- Multi-threading is not enabled by default in the data_engine probe. To increase data_engine performance, you can increase the number of threads by using the thread_count_insert parameter in Raw Configure. The optimum thread count is highly dependent on several factors, including:
 - The number of CPUs running on the system.
 - The number of RN tables in the UIM database.
 - The size of your UIM deployment.

Hub Bulk Size

- By default, the hub bulk size is set to 20. A low bulk size is optimal for small environments with small message rate throughput in the data_engine. However, if your UIM deployment has a high message rate you can increase the hub_bulk_size parameter in Raw Configure to values in the 1800-2000 range. This action increases the number of QoS messages that are sent at once between the hub and data_engine.

Chapter 4: Configuration Details

This section describes the configuration information and options available through the Admin Console data_engine configuration GUI. The navigation pane organizes data_engine configuration into the following nodes:

- [data_engine](#) (see page 23)
- [Scheduler](#) (see page 30)
- [Quality of Service](#) (see page 29)
- [Database Configuration](#) (see page 24)

To access the data_engine configuration interface, select the robot that the data_engine probe resides on in the Admin Console navigation pane. In the Probes list, click the arrow to the left of the probe and select Configure.

data_engine

Navigation: data_engine

This section lets you view probe and QoS information, change the log level, and set data management values.

Probe Information

This section provides the basic probe information and is read-only.

General Configuration

This section provides general configuration details.

- **Log Level:** Sets the amount of detail that is logged to the log file.
- **Data Management default values:** The default settings for data maintenance. These settings apply to all QoS settings unless they have been individually overwritten in the Quality of Service settings.

- **Compress data before delete:** If selected, then by default, data from the raw (RN) tables is summarized into the Hourly (HN) tables, and then deleted from the raw tables. In addition, the data from the Hourly (HN) tables is summarized into Daily (DN) tables, and then deleted from the Hourly tables. This is only completed before a delete is performed.
- **Delete raw data older than:** Raw data older than the indicated number of days is deleted.
- **Delete historic data older than:** Hourly table data older than the indicated number of days is deleted.
- **Delete daily average data older than:** Daily table data older than the indicated number of days is deleted.

Quality of Service Type Status

This section provides data regarding the QoS tables and is read-only.

Note: The status information is created based on statistics that are generated by the database provider. If incorrect information is displayed, you might need to update the table statistics. For more information, refer to [Update the Table Statistics](#) (see page 34).

Database Configuration

Important! The database connection properties should only be changed in limited circumstances such as recovery operations. Changing the Database Vendor can cause connection issues. If you are changing database vendors, CA recommends reinstalling CA Unified Infrastructure Management (UIM).

The Database Configuration section allows you to specify the database connection settings. These settings are different for each database vendor:

- [MySQL](#) (see page 25)
- [Microsoft](#) (see page 26)
- [Oracle](#) (see page 28)

To test the connection for all vendors, click **Actions, Test Connection** at the top of the screen.

MySQL

Navigation: data_engine>Database configuration>MySQL

This section lets you configure the connection options for a MySQL database.

- **Schema:** The database schema name.
- **Server Host:** The database server name or IP address.
- **Port:** The port number to connect to the database server.
Default: 3306
- **Username:** The login user name.
- **Password:** The login user password.
Restriction: The password does NOT contain any special characters (such as ";").

Microsoft SQL Server

Navigation: data_engine>Database configuration>Microsoft

This section lets you configure the connection and maintenance options for a Microsoft SQL Server database.

- **Provider:** The SQL server provider.
- **Initial Catalog:** The database name.
- **Data Source:** The database server.
- **User ID:** The login user.
- **Password:** The login user password.
Restriction: Password cannot contain any special characters (such as ";").
- **Parameters:** Other parameters for the OLEDB connection.
- **Partition Data Tables:** Select this check box to perform partitioning on the raw sample data tables.
Note: This option is only available for Microsoft SQL Server Enterprise Edition. It is not available in the Microsoft SQL Express edition.
- **Index Maintenance:** Perform table reindexing with other maintenance routines, which by default are executed every 24 hours.
- **Compression mode:** The method that is used for data compression:
 - **None:** No compression occurs.
 - **Page:** Optimizes storage of multiple rows in a page, a super-set of row compression.
 - **Row:** Stores fixed-length data types in variable-length storage format.
- **Maintenance mode:** How the indexes are maintained:
 - **Dynamic:** Maintenance is performed based on the index statistics.
 - **Reorganize:** Maintenance is performed using the "alter index ... reorganize" SQL Server script.
 - **Rebuild:** Maintenance is performed using the "alter index ... rebuild" SQL Server script.

Important! It is not possible to rebuild the index for single partitions prior to SQLServer 2014. You can only reorganize individual partitions. Performing automatic indexing for large tables from the data_engine is discouraged, as indexing may not complete in a reasonable amount of time.
- **Online mode:** The effect of maintenance on concurrent use of the QoS tables:
 - **Dynamic:** The maintenance is determined by the edition of SQL Server. If SQL Server is the Enterprise Edition, then Online mode is used for maintenance (if the chosen maintenance mode supports it); otherwise, Offline mode is used.

- **Online:** The QoS tables are available for update and query during the data maintenance period. Online mode offers greater concurrency but demands more resources.
- **Offline:** The QoS tables are unavailable for update and query during the data maintenance period.
- **Fragmentation level: Low threshold:** If the fragmentation for an index is less than the low threshold percent value, then no maintenance is performed.
- **Fragmentation level: High threshold:** If dynamic maintenance mode is selected and fragmentation is between the low and high threshold percentages, then the Reorganize mode is used; otherwise the Rebuild mode is used.

Note: This option is only available for Microsoft SQL Server Enterprise Edition. It is not available when using the Microsoft SQL Express edition.

- **Index name pattern:** The indexes that are maintained.

Default: Blank (a blank entry results in all indexes being considered for maintenance).

Oracle

Navigation: data_engine>Database configuration>Oracle

This section lets you configure the connection and database maintenance options for an Oracle database.

- **Hostname:** The database server name or IP address
- **Port:** The port number to connect to the database server
- **Username:** The login user name
- **Password:** The login user password.
Restriction: Password cannot contain any special characters (such as ";").
- **Service Name:** The Oracle SID or Service name.
- **Partition Data Tables:** Select this check box to perform partitioning on the raw sample data tables.
- **Index Maintenance:** Perform table re-indexing with other maintenance routines, which by default are executed every 24 hours.
- **Online mode:** The effect of maintenance on concurrent use of the QoS tables:
 - **Dynamic:** The maintenance is determined by the edition of Oracle. If Oracle is the Enterprise Edition, then Online mode is used for maintenance; otherwise, Offline mode is used.
 - **Online:** The QoS tables are available for update and query during the data maintenance period. Online mode offers greater concurrency but demands more resources.
 - **Offline:** The QoS tables are unavailable for update and query during the data maintenance period.
- **Fragmentation Level (%):** The percentage of fragmentation required before index maintenance is performed.

Quality of Service

Navigation: data_engine>Quality of Service

The Quality of Service section displays the attributes for the QoS metrics.

- **Name:** The QoS type name.
- **Description:** Description of the QoS type.
- **QoS Group:** The QoS group is a logical group to which the QoS belongs (optional).
- **Unit:** The unit of the QoS data (the abbreviated form of the QoS data unit).
- **Has Max Value:** The data type has an absolute maximum.
Example: disk size or memory usage
- **Is Boolean:** The data type is logical (yes/no).
Example: Host is available/unavailable or printer is up/down.
- **Type:** Different data types:
 - **0** = Automatic (The sample value is read at fixed intervals, which are set individually for each probe).
 - **1** = Asynchronous (The sample value is read only when the value changes, and the new value is read).
- **Override Raw Age:** Select this check box to override the raw age of the QoS metric.
- **Raw Age:** The number of days you want to retain the QoS metric information.
- **Override History Age:** Select this check box to override the history age for the QoS metric.
- **History Age:** The number of days you want to retain the history information.
- **Override Daily Average Age:** Select this check box to override the daily average age for the QoS metric.
- **Daily Average Age:** The number of days you want to retain the daily average information.
- **Override Compression:** Select this check box to override compression settings for data in RN and HN tables.
- **Compress:** Raw data is summarized and aggregated into Hourly (or historic) data before it is deleted from the RN tables. This Hourly data is then summarized and aggregated into Daily data before it is deleted from the HN tables.

Scheduler

Navigation: data_engine>Scheduler

This section allows you to schedule database maintenance.

- **Start time** - Select either **Now** or a specific date and time. Selecting now begins the new database maintenance schedule immediately.
- **Ending** - Select **Never**, **After** x occurrences, or **By** a specific date and time.
- **Recurring** - select one of the following occurrence patterns:
 - **Minutely**
 - **Hourly**
 - **Daily** (including a specific time)
 - **Weekly** (including a specific time and days of the week)
 - **Monthly** (including occurrence, calendar day, and specific time)
 - **Yearly** (including month and specific time)

Chapter 5: Troubleshooting

Troubleshooting topics:

- [Viewing the Log File](#) (see page 31)
- [Corrupted QoS Definition Values](#) (see page 32)
- [Out of date information in the Quality of Service Type Status Table](#) (see page 34)
- [Check Partitioning Jobs](#) (see page 38)

If your problem is not addressed here:

- Look for a solution or ask other users for help on the [CA UIM Community Forum](#).
- Contact [Support](#).

Send us feedback with the "rate this page" link below. We will strive to include a solution in the next release of this document.

Viewing the Log File

Advanced users may find it helpful to view the log file. To view the log file, click the data_engine probe and select **View Log**. You also can modify the log file settings so that it retains more data for troubleshooting.

Corrupted QoS Definition Values

Symptom:

I see corrupted QoS Definition values in QoS data within reports or dashboards, for example user-defined units such as degrees F or C, watts, etc. being labeled "variant". Customized QoS definition units have been incorrectly replaced (overridden) by data_engine version 7.85 (NM Server 5.60) or data_engine 7.86 (NM Server 5.61).

Solution:

Versions 7.87 through 7.90 of data_engine include a patch utility to recover and restore customized QoS definition units incorrectly replaced (overridden) by data_engine version 7.85 (NM Server 5.60) or data_engine 7.86 (NM Server 5.61). This patch utility uses a conditional override approach that corrects the issue.

Follow these steps:

1. From a command prompt, navigate to the directory that contains the utility:
(Unix) `cd <UIM install location>/probes/slm/data_engine/tools`
(Windows) `CD <UIM install location>\probes\slm\data_engine\tools`
2. Execute the tool in "report" mode (-r flag set) with the java version installed with UIM Server (formerly name NM Server):
(Unix) `../..../jre/jre1.6.0_24/bin/java -jar qos_def_unit_repair_kit.jar -r`
(Windows) `..\..\..\jre\jre1.6.0_24\bin\java -jar qos_def_unit_repair_kit.jar -r`
3. The patch utility scans the S_QOS_DEF_SKIP_UNIT table in the database and finds QoS Definitions that are suspected to be corrupt.

Important: The S_QOS_DEF_SKIP_UNIT table holds QoS definition values that should not override what is sent by a QoS probe. This table is pre-populated with three values:

- variant
- none
- user defined

If you have defined additional custom values, which have been incorrectly overridden by a previous data_engine version, add these values as new rows to the S_QOS_DEF_SKIP_UNIT table prior to step 5 below, so that the patch utility will find, report, and fix them as well. (Use standard database management tools to connect to the UIM database and add the rows and new values to the S_QOS_DEF_SKIP_UNIT table.)

The report generated by the utility shows corrupt QoS Definitions (if any) and the probe or probes associated with that data. Here is an excerpt from an example report:

Listing current problems...

```
S_QOS_DEFINITION { name=QOS_CPUSAMPLECOUNT, qosDefId=12,
group=QOS_VMWARE, unit=none,
unitShort=sc }
```

```
S_QOS_DATA { source=esxiqa1.i9.x, target=CPU sample count,
origin=w2k8-vm0hub,
host=10.0.0.1, robot=w2k8-vm0, probe=vmware }
```

```
S_QOS_DEFINITION { name=QOS_SNMP_VARIABLE, qosDefId=11,
group=QOS_SNMP_VARIABLE, unit=variant,
unitShort=value }
```

```
S_QOS_DATA { source=w2k8-vm0, target=interfaces.ifTable.ifREntry.ifOutOctets.1,
origin=w2k8-vm0hub, host=10.0.0.1, robot=w2k8-vm0, probe=snmpget }
```

```
S_QOS_DATA { source=w2k8-vm0, target=system.sysServices.0,
origin=w2k8-vm0hub,
host=10.0.0.1, robot=w2k8-vm0, probe=snmpget }
```

End of current problems list.

Running in report-only mode

4. Run the utility without the `-r` flag set to make the repairs:

```
(Unix) ../../../../jre/jre1.6.0_24/bin/java -jar qos_def_unit_repair_kit.jar
```

```
(Windows)..\..\..\jre\jre1.6.0_24\bin\java -jar qos_def_unit_repair_kit.jar
```

The utility subscribes to the NimBUS to receive QoS Definitions and corrects any in the database that have been incorrectly replaced.

5. Restart each of the probes listed in the report. In the example above the vmware and snmpget probes would be restarted so that the correct units are received. Probes are identified at the end of each `S_QOS_DATA` listing in the report. For example:

```
S_QOS_DATA { source=esxiqa1.i9.x, target=CPU sample count,
origin=w2k8-vm0hub,
host=10.0.0.1, robot=w2k8-vm0, probe=vmware }
```

Once the patch utility has repaired all QoS Definitions it does not need to be run again and no other action is required.

Out of date information in the Quality of Service Type Status Table

Symptom:

I see out of date information in the Quality of Service Type Status Table.

Solution:

Table statistics are collected by database software to create the best execution plan for a query. Some examples of collected statistics include:

- Rows that are stored in a table
- Available indexes
- How many pages store the rows

The data_engine probe uses these table statistics to generate the Quality of Service Type Status table. Table statistics can be manually updated if they become out of date.

The procedure for updating the table statistics is different for each database vendor:

- [Table Statistics for Oracle](#) (see page 34)
- [Table Statistics for MySQL](#) (see page 35)
- [Table Statistics for Microsoft SQL Server](#) (see page 36)

Important! Updating table statistics significantly impacts performance on all database platforms, especially on larger databases.

Table Statistics for Oracle

To receive correct statistics in Oracle, following the instructions found at:

- http://docs.oracle.com/cd/B28359_01/server.111/b28274/stats.htm#PFGRF003

Table Statistics for MySQL

To receive correct statistics in MySQL, run one of the following queries on all RN_, HN_, BN_, and DN_ tables:

- `ANALYZE table RN_QOS_DATA_XXXX ESTIMATE STATISTICS;` - Samples from the table are taken and stored in the data dictionary.
- `ANALYZE table RN_QOS_DATA_XXXX COMPUTE STATISTICS;` - The entire table is analyzed using a full table scan and stored in the data dictionary.

Using the ANALYZE command in MySQL can be a time-consuming operation, especially for large databases. Only perform the command sporadically and do not use Automated Maintenance Tasks.

In MySQL, the ANALYZE command holds a read lock on tables, which can negatively impact database performance.

For more information, refer to the following MySQL documentation:

- <http://dev.mysql.com/doc/refman/5.5/en/tables-table.html>
- <http://dev.mysql.com/doc/refman/5.6/en/analyze-table.html>

Table Statistics for Microsoft SQL Server

Normally table statistics are automatically managed by the SQL Server. However, this functionality does not work in some cases.

For example, if you are performing bulk inserts using the OLE DB FastLoad API, the statistics for data tables are not automatically updated. This can lead to poor performance and extra work for SQL Server.

The data_engine probe contains options that can automatically update statistics for Microsoft SQL Server. The code is in a stored procedure in the SLM database that is named 'spn_de_UpdateStatistics'.

Updating tables statistics in Microsoft SQL server causes queries to recompile. For more information, refer to the following Microsoft SQL documentation:

- <http://msdn.microsoft.com/en-us/library/ms187348.aspx>
- [http://msdn.microsoft.com/en-us/library/ms187348\(v=sql.100\).aspx](http://msdn.microsoft.com/en-us/library/ms187348(v=sql.100).aspx)

The behavior can be controlled with the following variables in the raw configure menu:

Key name	Default value	Type
statistics_age	24	Time in hours. This means when the stored procedure is called, statistics that are older than this number will be updated. This value is used by the stored procedure, not data_engine itself. If this number is set to 0 (zero), statistics will be disabled and not be run at all by the data_engine.
statistics_pattern	RN_QOS_DATA%	String pattern to which data tables will be updated.
statistics_loglevel	0	Numbers 0 to 5 used by the stored procedure when logging to tbnLogging.

statistics_time_pattern	<not set>	<p>The scheduling string that determines when to run statistics. If this key is empty or not set, the same schedule that is defined for data management will be used. This means the statistics will be run when data_engine has finished index maintenance and data management.</p> <p>If this value is specified to a different schedule, the statistics will be updated independently of when data management is scheduled.</p> <p>The string will be used by the calendar scheduling library, which is used by various UIM components. It supports RFC2445. See short example below.</p>
-------------------------	-----------	--

Some string examples that are copied from the library help file.

```

/*****
** nimCalCreate - Creates a handle to a nimCal structure
*****/
** PARAMETERS:
**   char   *pattern - RFC2445 'weekly' or 'dts'
**   char   *start   - startdate: yyyy-mm-dd [hh:mm:ss] || NULL
**                                     : weekly format 1 or 2
**
**   start = 'yyyy-mm-dd [hh:mm:ss]' will expect the 'pattern' to comply with RFC2445.
**   = NULL results in setting start to 'now'
**   e.g.
**       h = nimCalCreate("DTSTART:19970610T090000|RRULE:FREQ=YEARLY;COUNT=10",NULL);
**       h = nimCalCreate("DTSTART:19970610T090000|RRULE:FREQ=YEARLY;COUNT=10","2007-07-25");
**
**   pattern = 'weekly' handles two 'start' formats:
**       1. 0123,10:00,14:00 [,NOT] (old NAS format)
**       2. MO=12:00-14:00,15:30-17:00;TU=08:00-16:00 (new, allow 8-16)
**
**       h = nimCalCreate("weekly","012,10:00,14:00");
**       h = nimCalCreate("weekly","MO=12:00-14:00,15:30-17:00;TU=08:00-16:00");
**       h = nimCalCreate("dts","2007-08-20 08:00,2007-08-27 08:00,2007-09-03 08:00,2007-09-10"
**
** Note: Free the handle using nimCalFree.
*****/

```

You can also create a schedule in NAS and use the resulting string from there or use data_engine scheduler to create a string.

Check Partitioning Jobs

Use the following SQL statement to determine which partitioning jobs are running:

```
SELECT sqltext.TEXT,  
req.session_id,  
req.status,  
req.command,  
req.cpu_time,  
req.total_elapsed_time  
FROM sys.dm_exec_requests req  
CROSS APPLY sys.dm_exec_sql_text(sql_handle) AS sqltext;
```

If the results display multiple "CREATE nonclustered index" statements, you have more than one partitioning job running.

To stop a partitioning job:

```
KILL [session_id]
```