



InformaticaTM

Operation

Informatica MDM - Product 360

Version: 10.1

Table of Contents

1	Database Operation	9
1.1	Database update	10
1.1.1	Prerequisites on Oracle	10
1.2	Maintenance tasks	10
1.2.1	Indices and Statistics	10
1.2.1.1	Create Maintenance Plan for SQL Server	11
1.2.1.2	Create Maintenance Plan for Oracle	16
1.2.2	Backup and Recovery	20
1.3	Tuning.....	20
1.3.1	MS-SQL Server	20
1.3.1.1	Activity Monitor	20
1.3.1.2	Data Collection	21
1.3.1.3	Memory Settings	21
1.3.1.4	Transaction Logs and the Recovery Mode	23
1.3.1.5	TempDB handling	23
1.4	Trouble shooting.....	23
1.4.1	Reporting.....	23
1.4.1.1	Usage statistics	23
1.4.1.2	Cleanup all temporary reports	27
1.4.1.3	Separate storage for reporting tables.....	28
1.4.2	Create missing entries Item/Product/Variant/Structure-attributes in repository key language	28
1.4.2.1	Solution	29
1.4.2.2	Consequences	29
1.4.3	MS-SQL Server	29
1.4.3.1	"The query processor could not start the necessary thread resources for parallel query execution"	29
1.4.3.2	"Transaction (Process ID XXX) was deadlocked on lock resources with another process and has been chosen as the deadlock victim. Rerun the transaction"	30
1.4.4	Oracle.....	30
1.4.4.1	Import/Export of schemas	30
1.4.4.2	Connection Issues	30
1.4.4.3	Oracle JDBC driver OutOfMemoryError / huge memory consumption.....	31
1.4.4.4	Oracle JDBC driver OutOfMemoryError (Cache)	32

1.4.4.5 Oracle ReportStoreTemp performance	32
1.4.4.6 Automatic expiration of passwords for Oracle DB users.....	32
1.5 Diagnosing and tuning poorly performing Oracle queries	32
1.5.1 Fixing frequency histogram issues.....	38
1.6 Oracle Import/Export for MAIN, MASTER and SUPPLIER Schemas	40
1.6.1 Preparation	40
1.6.1.1 Directory Object	40
1.6.2 Export	41
1.6.2.1 User Roles	41
1.6.2.2 Execute export datapump	41
1.6.3 Import.....	41
1.6.3.1 User Roles	41
1.6.3.2 Tablespaces.....	42
1.6.3.3 Execute Import Datapump	43
2 Server Operation	45
2.1 Maintenance Tasks	45
2.2 Monitoring	45
2.3 Tuning.....	46
2.3.1 Memory Settings	46
2.3.2 Garbage Collection	46
2.3.2.1 GC activity log.....	46
2.3.2.2 Analyse GC activity log & optimize	47
2.3.3 Communication Framework	47
2.3.4 Mass-Data Parallel Processing	47
2.3.5 Media Asset Parallel Management	47
2.3.5.1 CreateZipManagement	47
2.3.5.2 MediaAssetProvider special MBean	49
2.3.5.3 HeilerClassic special MBeans	50
2.3.6 Database Connection Pool	50
2.3.7 Logging	51
2.3.8 Import.....	52
2.3.8.1 Preferences regarding CPU usage.....	52
2.3.8.2 Preferences regarding memory usage	53
2.4 Trouble Shooting	54

2.4.1	Anti-Virus Software affecting performance	54
2.4.2	Media Asset Provider (Classic Provider).....	54
2.4.2.1	No preview are visible.....	54
2.4.3	Obtain a Memory Dump (Heap Dump) from the Java Virtual Machine	55
2.4.4	Enhanced Logging for Assortment Content Calculation.....	55
2.4.5	Server and client in different time zones	56
2.4.6	Cluster.....	56
2.5	SNMP Monitoring	57
2.5.1	Overview	57
2.5.2	Default Configurations.....	58
2.5.2.1	Agent.....	58
2.5.2.2	Network Management Station	58
2.5.2.3	Managed Objects.....	59
2.5.2.4	OID Configuration	59
2.5.2.5	Generation of OID configuration	59
2.5.2.6	Server Remote Command	59
2.5.2.7	Trap.....	60
2.5.2.8	Enable SNMP	60
2.5.2.9	Run application.....	61
2.5.3	Modifying configurations.....	61
2.5.3.1	Change the agent's port	61
2.5.3.2	Network Management Station	61
2.5.3.3	Change OID Configuration	62
2.5.3.4	Advanced configuration	64
2.6	DB Connection pool	65
2.6.1	High performance connection pool configuration example	66
2.6.1.1	Hibernate.....	66
2.6.1.2	UDA	68
2.7	Server Job Maintenance	70
2.7.1	Configure the life time of a job	70
2.7.2	Syntax	70
2.7.2.1	Example	70
2.7.3	Standard Job List.....	70
2.8	Soft Delete Cleanup Job	73
2.8.1	Why is there a concept of soft delete and a cleanup job that triggers the logging in the DB?	73

2.8.2	General information.....	73
2.8.3	Configuration	74
2.8.4	Schedule.....	74
2.8.5	Be cautious with mass deletes in such a setup	75
2.8.5.1	How do other customers solve such a constellation?.....	75
2.8.5.2	Possible Workarounds	75
2.9	Data Quality Operation	76
2.9.1	Logging	76
2.9.1.1	Rule engine execution.....	76
2.9.1.2	Data provisioning	76
2.9.1.3	Performance analyzing.....	77
2.9.2	Performance tuning.....	77
2.9.2.1	Cleanup of Status Entries	80
2.9.3	Issue reporting guide	81
2.9.3.1	Configuration setup	81
2.9.3.2	Information about the execution context.....	81
2.9.3.3	Logging data.....	81
2.9.3.4	Performance.....	81
2.9.3.5	Crashes	82
2.9.4	Configure rules that are affected by Dictionary Synchronization in order to increase performance of IDQ rule execution	82
2.9.4.1	Situation/ motivation	82
2.9.4.2	Solution	82
2.10	Check integrity of attribute names in key language	84
2.10.1	MSSQL.....	84
2.10.2	Oracle.....	85
2.11	Network Recovery	85
2.12	Monitoring with Micrometer.....	86
2.12.1	General monitoring approach.....	86
2.12.2	List of Product 360 custom metrics.....	87
2.12.3	Sample for monitoring a Product 360 server with Micrometer, Prometheus and Grafana	87
3	Desktop Operation	89
3.1	Maintenance Tasks	90
3.2	Tuning.....	90

3.2.1	Performance optimization of the "Documents" view	90
3.2.1.1	Filter	91
3.2.2	Import	92
3.3	Trouble Shooting	93
3.3.1	User Interface is not responding any more (client freeze)	93
3.3.1.1	Find the problem on the client	93
3.3.1.2	Find the problem on the server	93
3.3.1.3	Connection timeouts between Server and Client	94
3.3.2	Rich Text Editor	94
3.3.2.1	Unable to use the clipboard actions "Cut", "Copy" and "Paste"	94
3.3.3	Workaround to prevent the "Invalid Certificate" error (Mozilla)	95
3.4	Product Paradim (EGD) limitations.....	96
4	Control Center Manual.....	96
4.1	Prerequisite	97
4.2	Function.....	97
4.3	Login	98
4.4	Dashboard	98
4.4.1	Widget creation using templates	98
4.4.2	Customize widgets	99
4.5	Monitoring	100
4.5.1	Application Server Node.....	101
4.5.1.1	Server tab	101
4.5.2	Server Nodes	102
4.5.3	Server Node Details.....	102
4.5.3.1	Clients.....	103
4.5.3.2	Log	104
4.5.3.3	Thread Dump.....	104
4.5.3.4	SNMP.....	105
4.6	Deployment.....	105
4.6.1	Installation	106
4.6.2	License.....	107
4.6.3	Configuration	107
4.7	Support.....	107
4.7.1	CSM	108

4.8	Additional Features.....	108
4.8.1	Opening multiple tabs	108
4.8.2	Telemetrie Actions	109
5	Audit Trail Operation	109
5.1	Backup to File.....	110
5.2	Restore from File	117
6	Log File Overview	122
6.1	Additional Logging of Dataquality	123
6.2	Additional Logging of Change Summary and Datagraph	123
6.3	Additional Logging of Media Asset Processing	124
7	How to create a thread dump.....	124
7.1	Control Center	125
7.2	Console Window.....	125
7.3	VisualVM (JMX Interface).....	125
7.4	Analyzing Thread Dumps.....	127
8	How to enable Java Management Extensions (JMX).....	129
8.1	Production environment	130
8.2	Download and Installation of JMX Tooling.....	130
8.2.1	Connect on productive system	131

The operation manual documents actions which need to be performed in order to keep Informatica MDM - Product 360 up and running in a stable way. It contains the backup strategy as well as tuning and troubleshooting hints.

In order to meet enterprise level requirements of our current and future customers we try to steadily improve the performance of all parts of the Product 360 application. Optimal performance comes not without a cost, you will need adjust the platform configuration to the actual use case and load scenarios you experience at the customer.

1 Database Operation

Combines maintenance, tuning and troubleshooting for the Product 360 Core Database server. Additionally to the recommendations of the operation guide, a well educated database administrator should be available for maintenance and tuning tasks.

- [Database update](#)(see page 10)
 - [Prerequisites on Oracle](#)(see page 10)
- [Maintenance tasks](#)(see page 10)
 - [Indices and Statistics](#)(see page 10)
 - [Create Maintenance Plan for SQL Server](#)(see page 11)
 - [Automatic execution \(MSSQL2012\)](#)(see page 11)
 - [Manual execution](#)(see page 15)
 - [Create Maintenance Plan for Oracle](#)(see page 16)
 - [Height-Balanced Histograms on NVARCHAR2 Columns](#)(see page 17)
 - [Extended Column Statistics](#)(see page 18)
 - [Backup and Recovery](#)(see page 20)
- [Tuning](#)(see page 20)
 - [MS-SQL Server](#)(see page 20)
 - [Activity Monitor](#)(see page 20)
 - [Data Collection](#)(see page 21)
 - [Memory Settings](#)(see page 21)
 - [Lock pages in memory](#)(see page 22)
 - [Transaction Logs and the Recovery Mode](#)(see page 23)
 - [TempDB handling](#)(see page 23)
- [Trouble shooting](#)(see page 23)
 - [Reporting](#)(see page 23)
 - [Usage statistics](#)(see page 23)
 - [Cleanup all temporary reports](#)(see page 27)
 - [Separate storage for reporting tables](#)(see page 28)
 - [Create missing entries Item/Product/Variant/Structure-attributes in repository key language](#)(see page 28)
 - [Solution](#)(see page 29)
 - [Consequences](#)(see page 29)
 - [MS-SQL Server](#)(see page 29)
 - ["The query processor could not start the necessary thread resources for parallel query execution"](#)(see page 29)
 - ["Transaction \(Process ID XXX\) was deadlocked on lock resources with another process and has been chosen as the deadlock victim. Rerun the transaction"](#)(see page 30)
 - [Oracle](#)(see page 30)
 - [Import/Export of schemas](#)(see page 30)
 - [Connection Issues](#)(see page 30)
 - [Oracle JDBC driver OutOfMemoryError / huge memory consumption](#)(see page 31)
 - [Oracle JDBC driver OutOfMemoryError \(Cache\)](#)(see page 32)
 - [Oracle ReportStoreTemp performance](#)(see page 32)
 - [Automatic expiration of passwords for Oracle DB users](#)(see page 32)

1.1 Database update

1.1.1 Prerequisites on Oracle

In case the Product 360 database is running in archive log mode (most likely), make sure to turn archive log mode off before executing the database update setup.

Otherwise you might risk exceeding available archive log disk space, in which case any database activity will be suspended immediately, and it will stay suspended until there is free available disk space again.

1.2 Maintenance tasks

1.2.1 Indices and Statistics

The database uses indices to increase the performance for read access on columns. In Product 360 we have indices on the columns which are typically used for selecting data, for example, all logical keys are at least part of an index. Additionally to that, the database uses statistics about the data in the tables in order to choose the most effective execution plan. Indices as well as statistics need to be maintained by the database administrator on a regular basis. Usually this is being done by creating a Maintenance Task which executes every day.

However, the typical maintenance task is more something of a bulldozer in maintaining. This means, it just rebuilds or reorganizes all indices no matter if needed or not and it also kills all statistics and rebuilds them no matter if needed or good for the statistic. Sometimes statistics which are very accurate get blindly rebuild and the result is a statistic which is not that accurate anymore, leading to a lesser performance than what you could have.

So, the general recommendation (from Microsoft Consulting) is like this (pseudo code):

```
if (number of pages in index < 1000)
    do nothing
else
    if (fragmentation < 5%)
        do nothing
    else if (fragmentation between 5% and < 30%)
        reorganize index + update its statistics
    else
        rebuild index
```

Additionally to that you should update statistics for all non-indexed columns and monitor index fragmentation to adjust the fill factor.

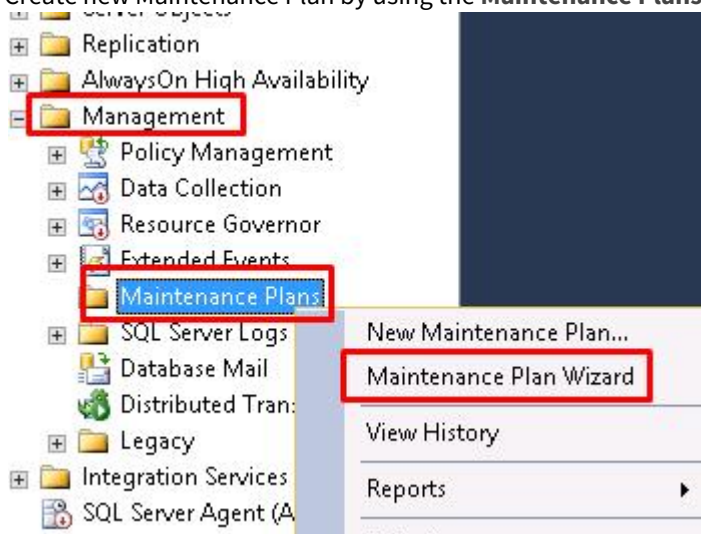
1.2.1.1 Create Maintenance Plan for SQL Server

All these steps can be combined in a single, easy to execute stored procedure which has generously been provided by Ola Hallengren. For a detailed description about the index and statistic maintenance script please refer to the online description which is available here:

<http://ola.hallengren.com/sql-server-index-and-statistics-maintenance.html>. It will create all necessary objects for you. (For your convenience we attached the script it also to this wiki page)

Automatic execution (MSSQL2012)

1. **Execute** the MaintenanceSolution.sql script on your database. Be sure that your **SQL Server Agent** (service) is running.
2. Create new Maintenance Plan by using the **Maintenance Plans Wizard** as shown below



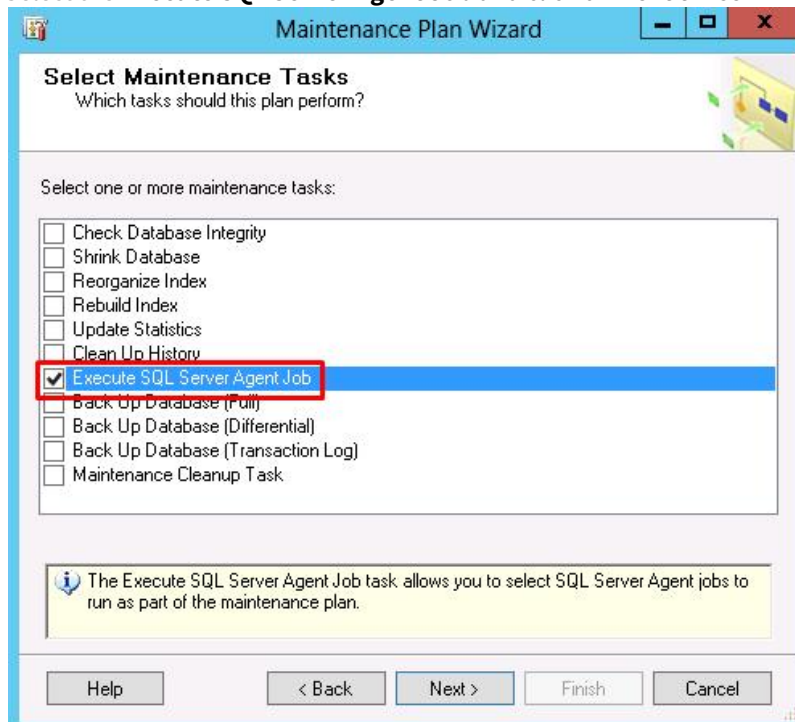
In case you get following error:




execute this statement:

```
sp_configure 'show advanced options', 1;
GO
RECONFIGURE;
GO
sp_configure 'Agent XPs', 1;
GO
RECONFIGURE
GO
```

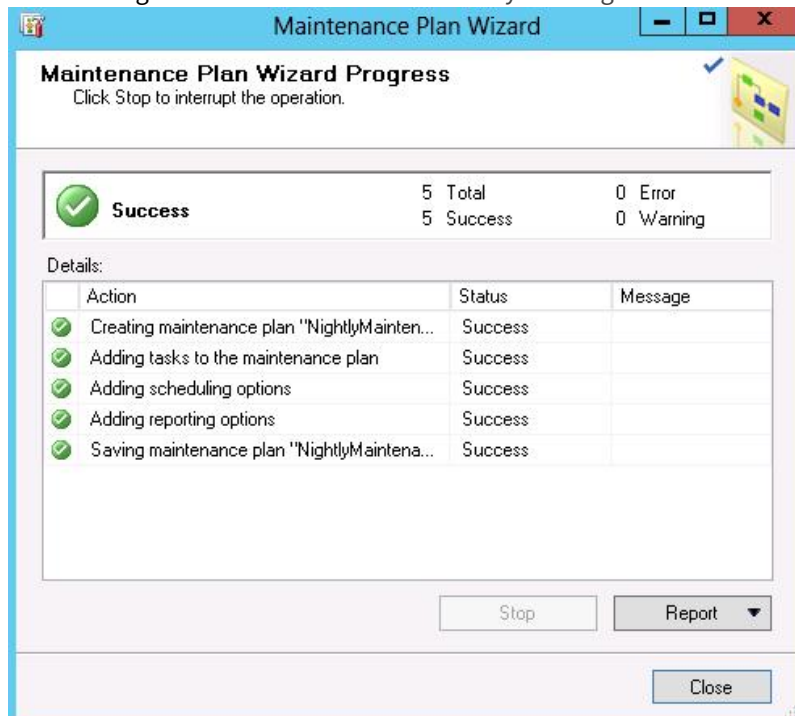
3. Type the name of your new maintenance plan and configure your job schedule, then click on **Next**
4. Select the **Execute SQL Server Agent Job** and click on **Next twice**



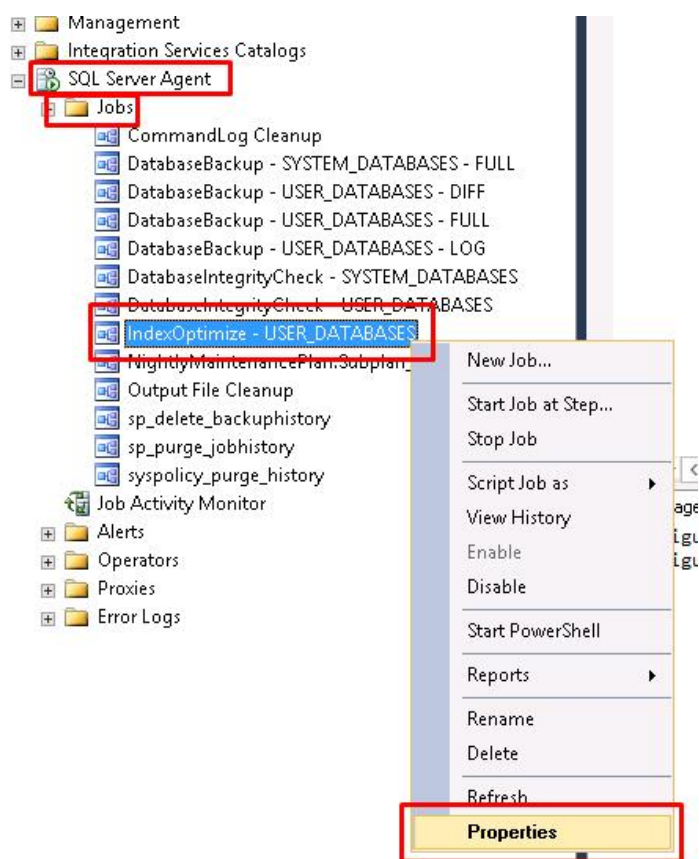
5. Choose **IndexOptimize** and click on **Next**



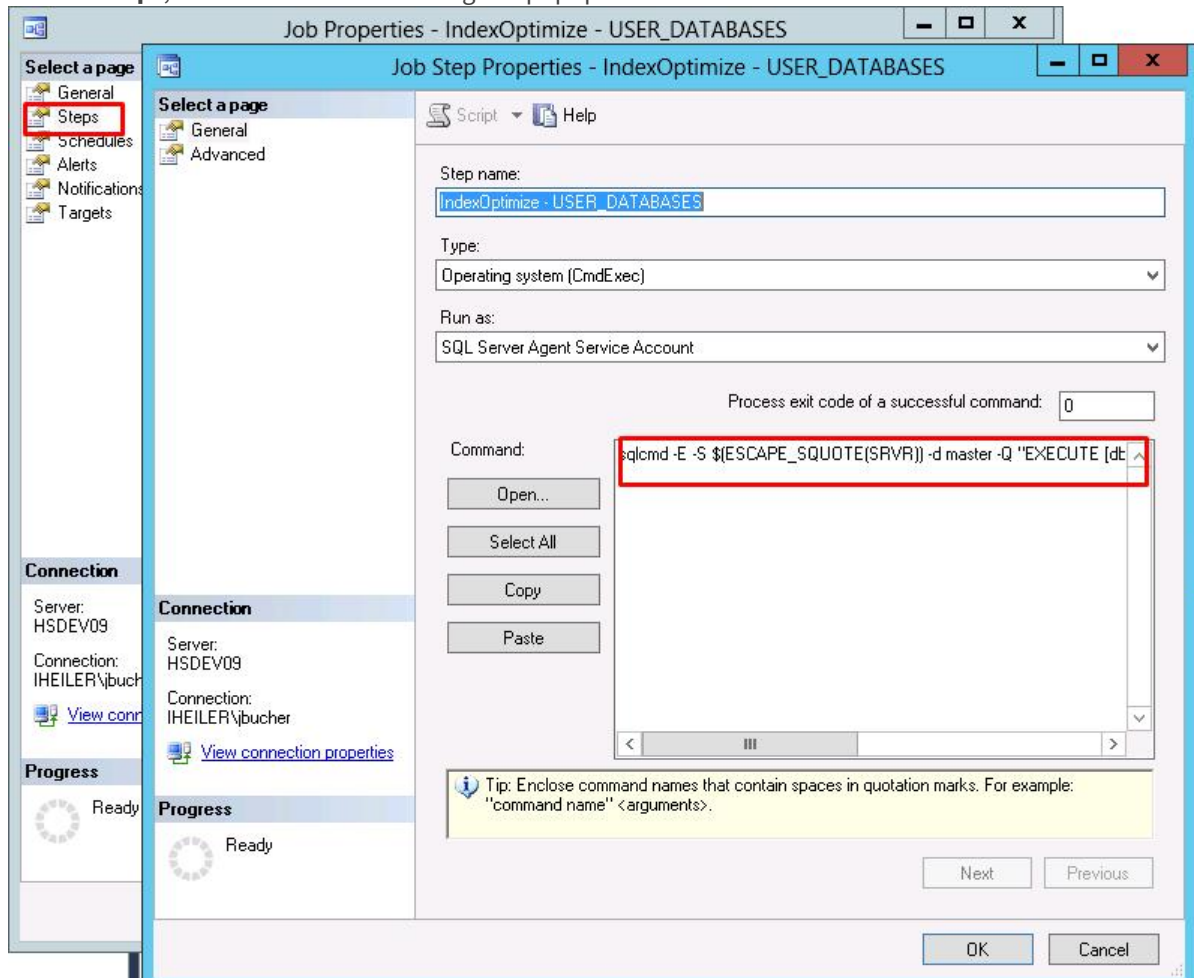
6. (Optional) Configure the report location and click on Next
7. After clicking **Finish** the task will check for any warnings or errors:



8. Click on the **SQL Server Agent** node from the left-hand side to expand the menu and open the **Jobs** node. Then right-click on **IndexOptimize** and go to the **Properties** view as shown below:



9. Click on **Steps**, then **Edit** and a new Dialog will popup.



10. **Modify** the command statement (in this example the job will be executed on all user databases)

```
sqlcmd -E -S $(ESCAPE_QUOTE(SRVR)) -d master -Q "EXECUTE [dbo].[IndexOptimize]
@Databases = 'USER_DATABASES', @FragmentationLow = NULL, @FragmentationMedium =
'INDEX_REORGANIZE,INDEX_REBUILD_ONLINE,INDEX_REBUILD_OFFLINE',
@FragmentationHigh = 'INDEX_REBUILD_ONLINE,INDEX_REBUILD_OFFLINE',
@FragmentationLevel1 = 5, @FragmentationLevel2 = 30, @UpdateStatistics =
'COLUMNS', @OnlyModifiedStatistics = 'Y', @LogToTable = 'Y'" -b
```

Make sure to reformat above statement to a valid format.

11. Click **OK**
12. Schedule the job in the **Properties** view

Manual execution

1. Download and **execute** the MaintenanceSolution.sql script.
2. **Execute** following SQL statement (in this example the index will get optimized on PIM_MAIN, PIM_MASTER and PIM_SUPPLIER)

SQL Statement to execute

```
EXECUTE dbo.IndexOptimize @Databases = 'PIM_MAIN, PIM_MASTER, PIM_SUPPLIER',
@FragmentationLow = NULL, @FragmentationMedium =
'INDEX_REORGANIZE,INDEX_REBUILD_ONLINE,INDEX_REBUILD_OFFLINE',
@FragmentationHigh = 'INDEX_REBUILD_ONLINE,INDEX_REBUILD_OFFLINE',
@FragmentationLevel1 = 5, @FragmentationLevel2 = 30, @UpdateStatistics =
'COLUMNS', @OnlyModifiedStatistics = 'Y'
```

Make sure to reformat above statement to a valid format.

1.2.1.2 Create Maintenance Plan for Oracle

If the Oracle database instance is fully dedicated to PIM, it is recommended to disable the system job which creates plan directives because these directives can cause performance issues.

Disable Automated Statistics Collection Job

```
BEGIN
  DBMS_AUTO_TASK_ADMIN.DISABLE
  (
    CLIENT_NAME => 'sql tuning advisor',
    OPERATION   => NULL,
    WINDOW_NAME => NULL
  );
END;
/
```

If the Oracle database instance is fully dedicated to PIM, it is recommended to disable the system job which automatically updates statistics on a daily basis because it does not create histograms:

Disable Automated Statistics Collection Job

```
BEGIN
  DBMS_AUTO_TASK_ADMIN.DISABLE
  (
    CLIENT_NAME => 'auto optimizer stats collection',
    OPERATION   => NULL,
    WINDOW_NAME => NULL
  );
END;
/
```

Instead, the DBA should create its own job for statistics maintenance. It should be executed on a periodic basis, recommended nightly but also ad-hoc after more than 20% of the data has been significantly changed (e.g. during a large import).

Here is an example script which covers default Product 360 installations, the DBA is encouraged to adjust it accordingly to the specific Product 360 environment:

Gather Statistics

```
--
*****
-- Statistics collection with SKEWONLY
-- * Automatically creates histograms on any column that shows a skew in data
distribution
-- * Despite time consumption it is always recommended to use ESTIMATE_PERCENT = 100
--
*****
BEGIN
  DBMS_STATS.GATHER_SCHEMA_STATS(
    OWNNAME      => 'PIM_MAIN',
    OPTIONS       => 'gather',
    ESTIMATE_PERCENT => 100,
    METHOD_OPT     => 'for all columns size skewonly',
    CASCADE       => true,
    DEGREE        => 8
  );
  DBMS_STATS.GATHER_SCHEMA_STATS(
    OWNNAME      => 'PIM_MASTER',
    OPTIONS       => 'gather',
    ESTIMATE_PERCENT => 100,
    METHOD_OPT     => 'for all columns size skewonly',
    CASCADE       => true,
    DEGREE        => 8
  );
  DBMS_STATS.GATHER_SCHEMA_STATS(
    OWNNAME      => 'PIM_SUPPLIER',
    OPTIONS       => 'gather',
    ESTIMATE_PERCENT => 100,
    METHOD_OPT     => 'for all columns size skewonly',
    CASCADE       => true,
    DEGREE        => 8
  );
END;
/
```

Height-Balanced Histograms on NVARCHAR2 Columns

In case there are more than 254 distinct values, Oracle can't use a frequency histogram anymore but switches to a height-balanced histogram instead.

There are known performance issues with height-balanced histograms on NVARCHAR2 columns and it is recommended to delete all of these height-balanced histograms (otherwise the optimizer might create bad execution plans due to unpopular / non-existent value choices).


Delete Height-Balanced Histograms for all NVARCHAR2 Columns

```

SET SERVEROUTPUT ON
BEGIN
  FOR COL_ITEM IN (
    SELECT stats.OWNER, stats.TABLE_NAME, stats.COLUMN_NAME
    FROM DBA_TAB_COL_STATISTICS stats
    INNER JOIN DBA_TAB_COLUMNS cols
    ON stats.TABLE_NAME = cols.TABLE_NAME
    AND stats.COLUMN_NAME = cols.COLUMN_NAME
    WHERE cols.DATA_TYPE = 'NVARCHAR2'
    AND stats.HISTOGRAM <> 'NONE'
    AND stats.NUM_DISTINCT > 254
    AND stats.OWNER = cols.OWNER
    AND stats.OWNER IN ('PIM_MAIN','PIM_MASTER','PIM_SUPPLIER')
    ORDER BY stats.OWNER, stats.TABLE_NAME, stats.COLUMN_NAME
  )
  LOOP
    DBMS_OUTPUT.PUT_LINE( 'Deleting height-balanced histogram for: ' ||
COL_ITEM.OWNER || '."' || COL_ITEM.TABLE_NAME || '."' || COL_ITEM.COLUMN_NAME ||
' "' );
    DBMS_STATS.DELETE_COLUMN_STATS (OWNNAME=>COL_ITEM.OWNER, TABNAME=>'"' ||
COL_ITEM.TABLE_NAME || '"', COLNAME=>'"' || COL_ITEM.COLUMN_NAME || '"',
COL_STAT_TYPE=>'HISTOGRAM');
  END LOOP;
END;
/

```

Extended Column Statistics

 In case you are running Oracle 12.1 be aware of the "extended column statistics" feature. We strongly recommend turning off this feature, and to make sure that there aren't any extended statistics columns in all PIM schemas.

The following script can be used to delete system generated extended statistics:

Delete Extended Column Statistics

```

SET SERVEROUTPUT ON
BEGIN
  FOR c IN
  (
    SELECT OWNER, TABLE_NAME, DBMS_LOB.SUBSTR(EXTENSION, 3000) X
    FROM DBA_STAT_EXTENSIONS
    WHERE OWNER LIKE 'PIM_%'
    AND CREATOR = 'SYSTEM'
    AND DROPPABLE = 'YES'
    ORDER BY TABLE_NAME, X
  )
  LOOP
    DBMS_OUTPUT.PUT_LINE('Deleting system generated column stats for "' ||
c.OWNER || '".' || c.TABLE_NAME || ' : ' || c.X);
    DBMS_STATS.DROP_EXTENDED_STATS(c.OWNER, '"' || c.TABLE_NAME || '"', c.X );
  END LOOP;
END;
/

```

It is recommended to disable the extended column statistics feature.

Disable Extended Column Statistics

```

ALTER SYSTEM SET "_optimizer_enable_extended_stats"=FALSE SCOPE=BOTH;

```

If they can't be disabled entirely, the following script can be used to disable extended column statistics only for all PIM schemas (possible starting with 12.2):

Disable Extended Column Statistics (>= 12.2)


```

BEGIN
  FOR TABLE_ITEM IN (
    SELECT t.OWNER, t.TABLE_NAME
    FROM DBA_TABLES t
    WHERE t.OWNER IN ('PIM_MAIN', 'PIM_MASTER', 'PIM_SUPPLIER')
    ORDER BY t.OWNER, t.TABLE_NAME
  )
  LOOP
    DBMS_STATS.SET_TABLE_PREFS(TABLE_ITEM.OWNER, '"' || TABLE_ITEM.TABLE_NAME || '"',
'AutoStatExtensions', 'OFF');
  END LOOP;
END;
/

```

1.2.2 Backup and Recovery

Usually the backup and recovery strategy is defined by the customers IT department and it's dba. Product 360 core works fine with nightly backups using the "simple" recovery mode of the database, as well as transaction log backups during the day using the "full" recovery mode. Choose whatever your customer required in order to feel safe.

 If you switch the schemas to recovery mode "full" you have to set up the transaction log backups. If you don't your transaction logs will keep on growing eventually without limits!

The backup plan should include the MASTER, SUPPLIER and MAIN schemas as well as all the runtime folders like import, export where the import and export files are stored.

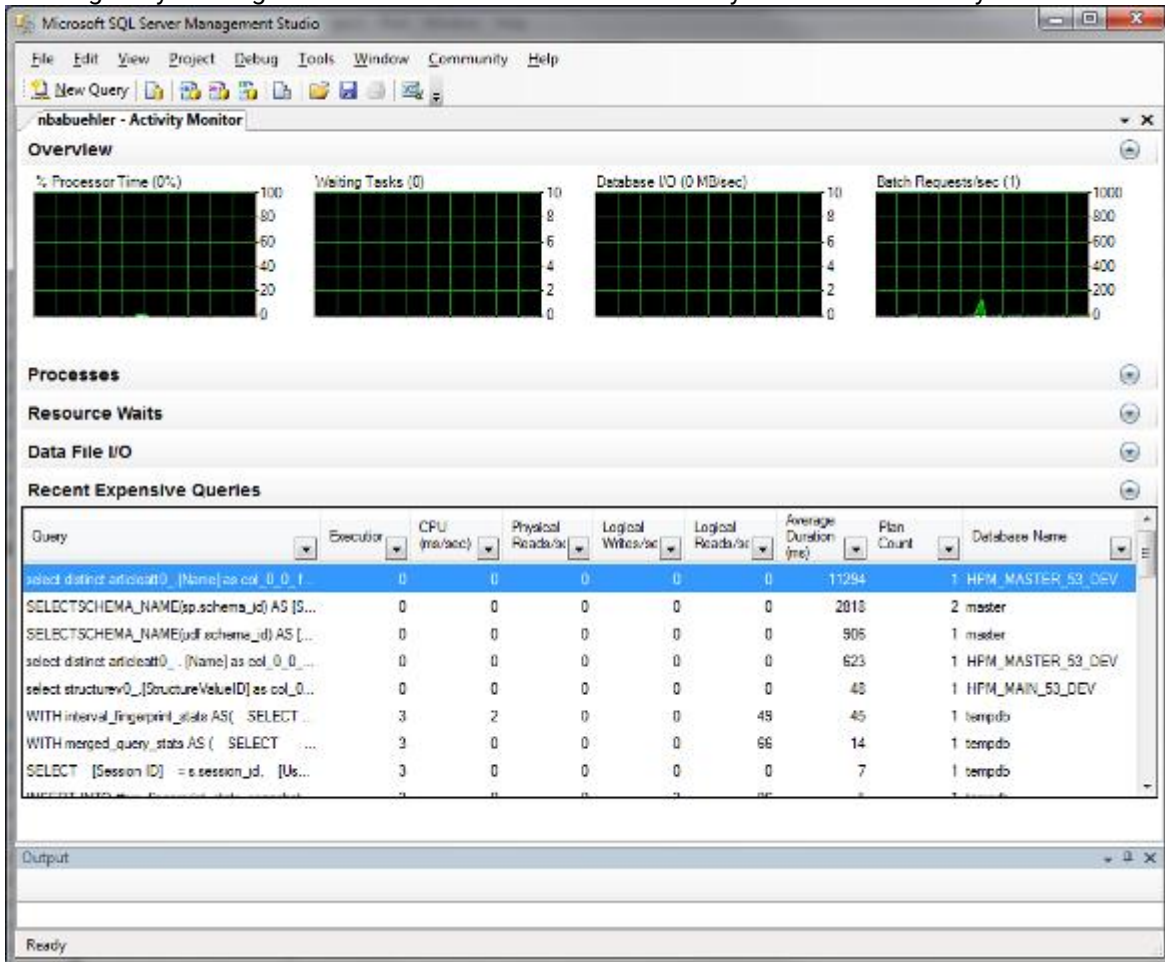
1.3 Tuning

1.3.1 MS-SQL Server

1.3.1.1 Activity Monitor

Use the Activity Monitor which is available in SQL Server to quickly identify slow statements and connection problems. Please see this knowledge base article which explains how to open the activity monitor, as well as

this blog entry which gives a brief overview on what the activity monitor can do for you



1.3.1.2 Data Collection

The activity monitor is a quite effective tool to get an overview of the current load on the database. However, sometime it's not that easy to produce the load on the application right at the time you're looking at the activity monitor. Sql Server 2008 has a new feature called "Data Collection". It might be a good idea to setup this feature and collect performance data for some days or weeks in order so see trends.

For a detailed description as how to setup this feature please refer to the MS-SQL Server Books Online: <http://msdn.microsoft.com/en-us/library/bb677356.aspx>

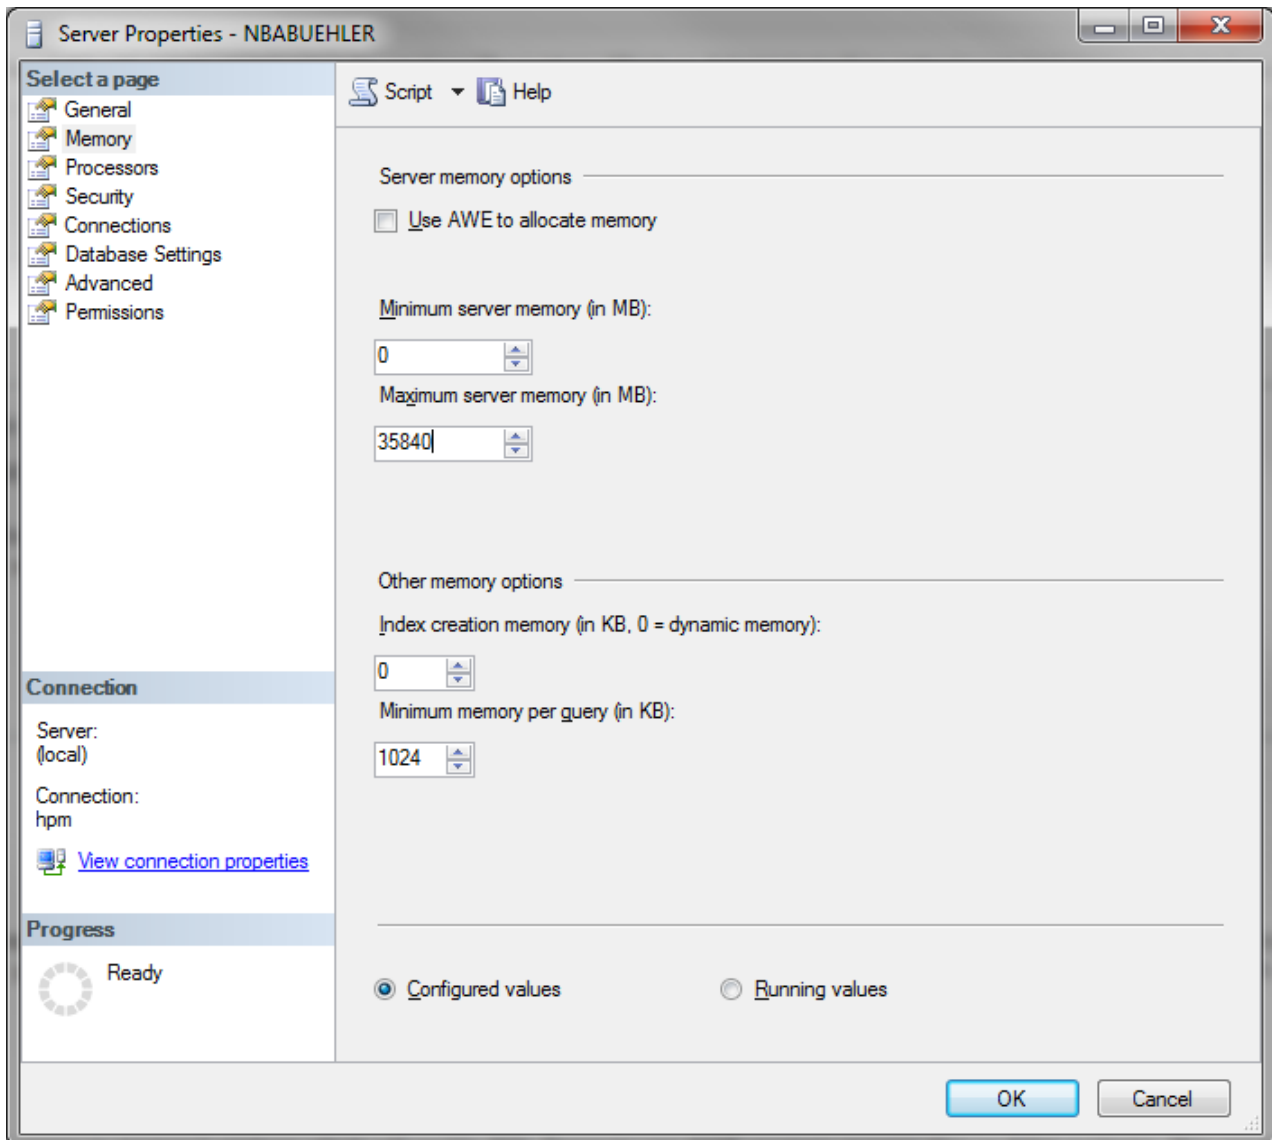
⚠ Please note that the data collection feature and the jobs it creates can not be deleted again, they can only be "disabled" so they will no longer collect any data.

1.3.1.3 Memory Settings

You should always limit the maximum memory usage of the SQL Server. In case you have a single instance on the machine, you can limit the memory to the maximum physical memory minus one GB for the operating system.

For example: The machine has 36 GB of memory, you would configure the SQL Server to use $35 \times 1024 = 35840$ MB memory, leaving the operating system 1024 MB - which should be enough, usually.

You can configure this using the Microsoft Sql Server Management Studio. Open the properties dialog of your database and adjust the memory settings like shown below.



⚠ The "Use AWE to allocate memory" will have no effect in 64bit installations, SQL Server uses AWE access automatically when the user has "Lock pages in memory" privilege.

Lock pages in memory


Enable the SQL Server Service user with the "Lock pages in memory" privilege which will enable SQL Server to use AWE memory access also on 64bit systems. This results in a better performance for the memory access since SQL Server can directly access the memory and must not ask for it using the Virtual Memory Manager (VMM).

See also "Configuring the Server for Optimal Performance" (SQL Server 2008 Administration from Wrox Press, Chapter 11, Page 434)

1.3.1.4 Transaction Logs and the Recovery Mode

In a productive environment you should define the initial size of the database log files to the expected maximum.

For both data and log files we recommend setting a proper growth value because the default of 64MB is too small.

 A data base growth action always "stops the world" of the database until the files are enlarged. When this will occur often it will be a serious performance problem.

In case you are operating the Product 360 databases in **FULL** recovery mode please make sure that you take backups of the transaction logs on a regular basis.

Otherwise the transaction logs will grow endlessly because SQL Server will not re-use the existing space until a backup was made.

1.3.1.5 TempDB handling

The TempDB will be used a lot since Product 360 uses transaction mode **READ_COMMITTED_SNAPSHOT** starting with PIM 7.1.

We recommend to tune the storage and the amount of data files according to the database load.

1.4 Trouble shooting

1.4.1 Reporting

A key functionality of the Product 360 - Server is the so called "reporting". A report in our sense is more or less only a list of internal database ID's. The report therefore reflects an arbitrary amount of "objects" in Product 360 - Server, identified by their ID. Reporting is a crucial part of Product 360 - Server since many functions rely on reports. For example the table based views of root entities like items or products, the export and the delete operations as well.

1.4.1.1 Usage statistics

The Informatica Software Support might ask you to provide usage statistics for the reporting in case we need to trouble shoot an issue of yours, for these you need to execute the following script on each data source (main, master and supplier):

Usage Statistic MSSQL

```

1  CREATE TABLE [#ReportInstances]
2  (
3      [ReportInstance] NVARCHAR(30)
4  )
5  INSERT INTO [#ReportInstances] ( [ReportInstance] ) VALUES (N'ReportStore')
6  , (N'ReportStoreTempA'), (N'ReportStoreTempB')
7
8  DECLARE @REPORT_TABLE_NAME NVARCHAR(30)
9  DECLARE @ReportInstance NVARCHAR(30)
10 DECLARE @SQL_STRING NVARCHAR(MAX)
11 DECLARE @TMP_SQL NVARCHAR(MAX)
12
13 DECLARE C0 CURSOR FOR
14 SELECT [ReportInstance]
15 FROM [#ReportInstances]
16 ORDER BY [ReportInstance]
17 OPEN C0
18 FETCH NEXT FROM C0 INTO @ReportInstance
19 WHILE @@FETCH_STATUS = 0
20 BEGIN
21     SET @SQL_STRING = ''
22     DECLARE C1 CURSOR FOR
23     SELECT TABLE_NAME
24     FROM INFORMATION_SCHEMA.TABLES
25     WHERE TABLE_NAME LIKE @ReportInstance + '[0-9]%'
26     ORDER BY CAST(REPLACE(TABLE_NAME, @ReportInstance, '') AS INT)
27     OPEN C1
28     FETCH NEXT FROM C1 INTO @REPORT_TABLE_NAME
29     WHILE @@FETCH_STATUS = 0
30     BEGIN
31         IF LEN(@SQL_STRING) > 0
32         BEGIN
33             SET @SQL_STRING = CONCAT(@SQL_STRING, ' UNION ALL ')
34         END
35         SET @SQL_STRING = CONCAT(@SQL_STRING, 'SELECT * FROM dbo.[' +
36 @REPORT_TABLE_NAME + ']' )
37         FETCH NEXT FROM C1 INTO @REPORT_TABLE_NAME
38     END
39     CLOSE C1
40     DEALLOCATE C1
41
42     SET @TMP_SQL = CONCAT('SELECT COUNT(*) AS [Number of Report IDs from
43 ', @ReportInstance, '] FROM (', @SQL_STRING, ') X')
44     PRINT @TMP_SQL
45     EXECUTE sp_executesql @TMP_SQL
46
47     SET @TMP_SQL = CONCAT('SELECT COUNT(*) AS [Number of Temporary Report
48 IDs from ', @ReportInstance, '] FROM (', @SQL_STRING, ') X WHERE ReportID
49 IN (SELECT ID FROM Report WHERE Purpose = 1)')

```



```

45      PRINT @TMP_SQL
46      EXECUTE sp_executesql @TMP_SQL
47
48      SET @TMP_SQL = CONCAT('SELECT COUNT(*) AS [Number of Permanent Report
IDs from ', @ReportInstance, '] FROM (', @SQL_STRING, ') X WHERE ReportID
IN (SELECT ID FROM Report WHERE Purpose = 3)')
49      PRINT @TMP_SQL
50      EXECUTE sp_executesql @TMP_SQL
51
52      SET @TMP_SQL = CONCAT('SELECT COUNT(*) AS [Number of Orphan Report IDs
from ', @ReportInstance, '] FROM (', @SQL_STRING, ') X WHERE ReportID NOT
IN (SELECT ID FROM Report)')
53      PRINT @TMP_SQL
54      EXECUTE sp_executesql @TMP_SQL
55
56      FETCH NEXT FROM C0 INTO @ReportInstance
57  END
58  CLOSE C0
59  DEALLOCATE C0
60
61  DROP TABLE [#ReportInstances]
62
63  SELECT COUNT([ID]) AS repTotalCount, Purpose, SUM([Count]) totalIdsCount
FROM [dbo].[Report] GROUP BY Purpose
64
65  SELECT count([ID]) AS numOfReports, [Count] AS repSize, COUNT([ID])*[Count]
AS totalSize FROM [dbo].[Report] WHERE Purpose=1 GROUP BY [Count] ORDER
BY totalSize DESC
66  SELECT count([ID]) AS numOfReports, [Count] AS repSize, COUNT([ID])*[Count]
AS totalSize FROM [dbo].[Report] WHERE Purpose=1 GROUP BY [Count] ORDER
BY numOfReports DESC
67  SELECT count([ID]) AS numOfReports, [Count] AS repSize, COUNT([ID])*[Count]
AS totalSize FROM [dbo].[Report] WHERE Purpose=1 GROUP BY [Count] ORDER
BY repSize DESC
68
69  SELECT [Type], SUM([Count]) AS totalSize FROM [dbo].[Report] WHERE
Purpose=1 AND [Count] < 10 GROUP BY [Type]
70  SELECT [Type], SUM([Count]) AS totalSize FROM [dbo].[Report] WHERE
Purpose=1 AND [Count] < 100 GROUP BY [Type]
71  SELECT [Type], SUM([Count]) AS totalSize FROM [dbo].[Report] WHERE
Purpose=1 AND [Count] < 200 GROUP BY [Type]
72
73  SELECT TOP 100 [ID] ,[Type] ,[Purpose] ,[Count] ,[CreationDate] FROM
[dbo].[Report] WHERE Purpose=1 ORDER BY CreationDate DESC

```

Usage Statistic Oracle

```

1  SET SERVEROUTPUT ON
2  SET LINESIZE 2000
3
4  DECLARE
5      REPORT_TABLE_NAME VARCHAR2(30);
6      SQL_STRING VARCHAR2(2000);
7      TMP_SQL VARCHAR2(2000);
8      TMP_RESULT NUMBER;
9      TEMP_NAME NVARCHAR2(30);
10     TYPE TEMP_ARRAY IS VARRAY(3) OF VARCHAR2(30);
11     ARRAY_TEMP_ARRAY := TEMP_ARRAY('ReportStore', 'ReportStoreTempA',
    'ReportStoreTempB');
12
13     CURSOR C1(REPORT_INSTANCE VARCHAR2)
14         IS SELECT TABLE_NAME
15            FROM USER_TABLES
16           WHERE REGEXP_LIKE(TABLE_NAME, '^' || REPORT_INSTANCE ||
    '\d')
17           ORDER BY CAST(REPLACE(TABLE_NAME, REPORT_INSTANCE) AS
    NUMBER);
18 BEGIN
19     FOR i IN 1..ARRAY_TEMP_ARRAY.COUNT LOOP
20         TEMP_NAME := ARRAY_TEMP_ARRAY(i);
21         SQL_STRING := '';
22
23         OPEN C1(TEMP_NAME);
24         LOOP
25             FETCH C1 INTO REPORT_TABLE_NAME;
26             EXIT WHEN C1%NOTFOUND;
27
28             IF LENGTH(SQL_STRING) > 0 THEN
29                 BEGIN
30                     SQL_STRING := SQL_STRING || ' UNION ALL ';
31                 END;
32             END IF;
33
34             SQL_STRING := SQL_STRING || ' SELECT * FROM "' || REPORT_TABLE_NAME
    || '"';
35         END LOOP;
36         CLOSE C1;
37
38         TMP_SQL := 'SELECT COUNT(*) AS "numIds' || TEMP_NAME || '" FROM (' ||
    SQL_STRING || ') X';
39         DBMS_OUTPUT.PUT_LINE(TMP_SQL);
40         EXECUTE IMMEDIATE(TMP_SQL) INTO TMP_RESULT;
41         DBMS_OUTPUT.PUT_LINE(TMP_RESULT || CHR(10));
42

```

```

43      TMP_SQL := 'SELECT COUNT(*) AS "numTempIds' || TEMP_NAME || '" FROM ('
      || SQL_STRING || ') X WHERE "ReportID" IN (SELECT "ID" FROM "Report"
WHERE "Purpose" = 1)';
44      DBMS_OUTPUT.PUT_LINE(TMP_SQL);
45      EXECUTE IMMEDIATE(TMP_SQL) INTO TMP_RESULT;
46      DBMS_OUTPUT.PUT_LINE(TMP_RESULT || CHR(10));
47
48      TMP_SQL := 'SELECT COUNT(*) AS "numPermIds' || TEMP_NAME || '" FROM ('
      || SQL_STRING || ') X WHERE "ReportID" IN (SELECT "ID" FROM "Report"
WHERE "Purpose" = 3)';
49      DBMS_OUTPUT.PUT_LINE(TMP_SQL);
50      EXECUTE IMMEDIATE(TMP_SQL) INTO TMP_RESULT;
51      DBMS_OUTPUT.PUT_LINE(TMP_RESULT || CHR(10));
52
53      TMP_SQL := 'SELECT COUNT(*) AS "numOprhanIds' || TEMP_NAME || '" FROM
(' || SQL_STRING || ') X WHERE "ReportID" NOT IN (SELECT "ID" FROM
"Report")';
54      DBMS_OUTPUT.PUT_LINE(TMP_SQL);
55      EXECUTE IMMEDIATE(TMP_SQL) INTO TMP_RESULT;
56      DBMS_OUTPUT.PUT_LINE(TMP_RESULT || CHR(10));
57
58  END LOOP;
59  END;
60  /
61
62  SELECT COUNT ("ID") AS "repTotalCount", "Purpose", SUM ("Count") AS
"totalIdsCount" FROM "Report" GROUP BY "Purpose";
63  SELECT COUNT ("ID") AS "numOfReports", "Count" AS "repSize", COUNT ("ID")*"
Count" AS "totalSize" FROM "Report" WHERE "Purpose"=1 GROUP BY "Count"
ORDER BY "totalSize" DESC;
64  SELECT COUNT ("ID") AS "numOfReports", "Count" AS "repSize", COUNT ("ID")*"
Count" AS "totalSize" FROM "Report" WHERE "Purpose"=1 GROUP BY "Count"
ORDER BY "repSize" DESC;
65
66  SELECT "Type", SUM ("Count") AS "totalSize" FROM "Report" WHERE "Purpose"=
1 AND "Count" < 10 GROUP BY "Type";
67  SELECT "Type", SUM ("Count") AS "totalSize" FROM "Report" WHERE "Purpose"=
1 AND "Count" < 100 GROUP BY "Type";
68
69  SELECT "Type", SUM ("Count") AS "totalSize" FROM "Report" WHERE "Purpose"=
1 AND "Count" < 200 GROUP BY "Type";
70  SELECT "ID", "Type", "Purpose", "Count", "CreationDate" FROM "Report"
WHERE "Purpose" = 1 ORDER BY "CreationDate" DESC;

```

1.4.1.2 Cleanup all temporary reports

The Product Manager uses an automatically scheduled system job for to delete temporary reports which are older than one day. This job is scheduled every hour (by default), therefore the Report table should not have any entry with Purpose=1 (temporary) which are older than 24 hours. However, it occurred in the past that this job could not fulfill its task in a reasonable amount of time or due some misconfiguration.

Prior to Version 6.0 this job deleted all reports (and their corresponding entries in the ReportStoreTemp table) in a **single transaction**. This may lead to an "explosion" of the transaction log since the ReportStoreTemp table might contain millions of items. In case you find old reports in the Report table (in any Product 360 - Server schema) you can execute the following SQL statements to completely delete all temporary reports in one small step, keeping the transaction log usage low.

Please note that you should shutdown the Product 360 - Server during this - since currently running operations might rely on a temporary report.

Delete ALL Temporary Reports	
1	DELETE FROM [Report] WHERE [Purpose] = 1
2	TRUNCATE TABLE [ReportStoreTemp]

With PIM - Server versions ≥ 6.0 the delete job is more aware of this problem. Big reports will be deleted one by one, small reports in a bulk delete operation. You can configure the boundaries of "big reports" in the `plugin_customization.ini` file. Additionally to that, the delete job can also be configured to delete reports which are younger than a day - for example a few hours.

1.4.1.3 Separate storage for reporting tables

On heavy load databases we recommend to put the corresponding reporting tables for each schema:

- Report
- ReportStore
- ReportStoreTempA
- ReportStoreTempB

on a high-performance storage separated from the rest of the Product 360 tables.

1.4.2 Create missing entries Item/Product/Variant/Structure-attributes in repository key language

This can happen if the language of the repository got changed during maintenance, if there is any custom code in the system which adds entries by not honoring the key language or an import which is ignoring the fact the the key language should be filled always. The repository language is configured in the `server.properties` and should get never touched in a running system.

Since PIM 7.0.01 the Product 360 Core will check this behavior during server startup. The server won't start if there are not all attributes maintained in the repository key language.


The occurred message in the server log file will look like this:

```
The database 'HPM_MAIN' contains invalid structure feature entries. '3' entries have
been found which aren't maintained in the repository language German.
```

1.4.2.1 Solution

Create the entries for the missing language. The provided script is checking the item attributes and structure attributes. It will lookup of the missing data entries and copy an existing entry with the lowest language ID for the missing repository language.

Because the name of the created attribute has to be unique, the name will be '<name of the attribute to copy><defined delimiter in the cmd file><id of the attribute to copy>'. This name will be cut at the cmd defined length (by default DB 'Name' column length).

 It is possible that this script creates invalid data if the attribute name of the the attribute to copy is as long as the defined max length in the cmd file and this name is already used by another attribute in the same context.

createMissingAttributes_mssql_pim7.rar

createMissingAttributes_oracle_pim7.rar

createMissingAttributes_mssql_pim8.rar

createMissingAttributes_oracle_pim8.rar

1.4.2.2 Consequences

The repository language **MUST NOT** be changed as soon as entity data such as items/products/variants or structures/structure groups have been created and exist in the database. In such a situation, the stability of the system can no longer be guaranteed since logical key fields most likely will contain null values.

1.4.3 MS-SQL Server

1.4.3.1 "The query processor could not start the necessary thread resources for parallel query execution"

During high-end benchmark tests the following error message has been thrown by the SQL Server "The query processor could not start the necessary thread resources for parallel query execution". After some research we came to the conclusion that the only currently known solution to this problem is to reduce the load on the database server. Several articles from Microsoft or Microsoft close resources explained this to be the only way to avoid this problem. See also this article from Microsoft: <http://msdn.microsoft.com/en-us/library/aa337448.aspx>

So, how to reduce the load on the database server? First approach is to reduce the number of concurrent connections using the [DB Connection pool](#)(see page 65) settings as described in the Tuning Advisory. Additionally to that please always inform the Informatica Software Support. It will be very helpful if you can describe which action lead to the error, steps to reproduce would be great too since we're constantly investigating this issue and try to find a better solution.

In case this problem occurs often on your system you might want to observe the processor queue, this can give us hints on the processor utilisation at the time the error occurred.

1.4.3.2 "Transaction (Process ID XXX) was deadlocked on lock resources with another process and has been chosen as the deadlock victim. Rerun the transaction"

In case you encounter database deadlocks you can use SQL Server Profiler to find the cause for the particular deadlock.

You need to run a server-side trace to capture deadlock information, you can use the script available here: <http://www.sqlservercentral.com/scripts/deadlock/66808/>

Once the deadlock has occurred, you can read the trace file with the script available here: <http://www.sqlservercentral.com/scripts/deadlock/66809/>

The content of the trace file will show which transactions resp. which SQL statements are responsible for the deadlock so further actions can be taken.

1.4.4 Oracle

1.4.4.1 Import/Export of schemas

To obtain some prepared notes on how to import export schemas from/to an Oracle instance, please raise a request with Informatica.

1.4.4.2 Connection Issues

If your project is having sporadic connection issues with Oracle database maybe these topics are helpful:

1. Maximum amount of processes. Because there were more than one application on the Oracle database the number of 300 was too small. We had good experiences with 1000 (minimum 300 x count of applications)
2. open_cursors, also too small values. We had good experiences with 1000 (minimum 300 x count of applications)
3. Log-Files are full. Mostly the Listener log file.

Switch off listener logging

Carry out following steps to configure the listener logging:

- a. Open [listener.ora](#) in `..\product\11.2.0\dbhome_1\NETWORK\ADMIN`
- b. Add **LOGGING_<Listernname> = OFF**
- c. Replace **<Listernname>** through your **listernname** e.q. 11GR2
- d. Save the changes
- e. Restart db and listener service.

Possible Values:

[ON| OFF]

4. tempdb full. Index rebuild after installation of standard classification systems!
5. Audit-Trail: waste logs also and makes them full. Towards description of Oracle the parameter „audit_trail“ is NOT deactivated per default! We deactivated this parameter on all of our instances.

⚠ Attention: when applying changes via enterprise manager console please ensure that the changes are applied in spfile, too.

1.4.4.3 Oracle JDBC driver OutOfMemoryError / huge memory consumption

Oracle JDBC driver may course OutOfMemoryError. The reason is that Oracle driver tries to allocate huge buffer for the incoming data. Sizes of these buffers depend on a single result set row size and 'prefetchRowSize' parameter of the prepared statement. By default this parameter is set to 10, but to increase list model performance it has been set to 10000 for all fragment SPs calls.

There are following different ways to avoid OutOfMemoryError in the oracle jdbc driver:

- Use jdbc driver parameter to limit buffer size in the UDA configuration file:

```
<property name="connectionProperties">oracle.jdbc.maxCachedBufferSize=<sizeInBytes>;</property>
```

- Set 'rowPrefetchSize' in the com.heiler.ppm.fragment.server.fragment extension point to override the global value for a certain fragment
- Set 'defaultRowPrefetchSize' in the com.heiler.ppm.fragment.server/preferences.ini or in the plugin_customization.ini

```
# Global value to control number of rows that will be fetch from the database in one
# round-trip
# using fragment stored procedure. It helps jdbc driver to find trade-off between
# memory consumption
# and db fetch performance. Default value set by many jdbc drivers is 10, but to
# improve list model data access
# we need to set this value to a much higher value.
#
# Some JDBC driver implementations use this value to estimate incoming raw data
# buffer size as
# defaultRowPrefetchSize*rowSize. Such buffers may have very large size (several
# hunderd megabyte).
# Set this parameter to a low value if application server has small amount
# of memory. You can also experiment with this value if garbage collection pessure is
# too high.
#
# This global value can be overridden per fragment using optional 'rowPrefetchSize'
# attribute of the
# com.heiler.ppm.article.core.server.Fragment extension point. It is recommended to
# use this attribute if
# the average number of rows returned by the fragment can be estimated in development
# time.
# defaultRowPrefetchSize = 10000
```

See Oracle JDBC MemoryManagement for more info

1.4.4.4 Oracle JDBC driver OutOfMemoryError (Cache)

The listModel (Product 360 - Server Table) access with NCLOBs may result in an OutOfMemoryError.

The reason of this dump is oracle driver cache for CLOB fields. JDBC statements hold references to huge byte arrays. This strategy usually reduces number of DB round-trips and improves performance but it takes a lot of memory.

Try setting **oracle.jdbc.FreeMemoryOnEnterImplicitCache** connection property to true. This will instruct Oracle driver to do cache clean-up more frequently.

see oracle driver doc for more detail: http://download.oracle.com/docs/cd/E14072_01/appdev.112/e13995/oracle/jdbc/pool/OracleDataSource.html

1.4.4.5 Oracle ReportStoreTemp performance

Product 360 - Server reporting framework uses ReportStore and ReportStoreTemp tables to save report ids. Under heavy load report save rate can be very high and ReportStoreTemp tables may grow up to dozens of million records.

If possible we recommend putting these tables on a high performance storage.

1.4.4.6 Automatic expiration of passwords for Oracle DB users

When setting up the Oracle DB and the users needed for Product 360 it is important to not set the password of each of them to automatically expire. Once the expiration date is met the Product 360 Server and components will not be able to start up anymore.

1.5 Diagnosing and tuning poorly performing Oracle queries

This is a big topic, but we'll focus in on what a bad query plan looks like, what a good query plan looks like, and one example of how to improve a poor performing query. There is no generic way to tune all poorly performing queries, but hopefully this page can educate you on the tools you need to identify and diagnose a poor performing query and an example of how this particular query was tuned. The first part of this how-to is PIM specific, but the important parts are only specific to Oracle.

The first step is to find the poor performing query. In our example, a PIM Desktop user issued a query and noticed poor performance:



The first step is to go through the logs to try to find the query. As PIM uses hibernate, the following categories (`org.hibernate.SQL` and `org.hibernate.type.descriptor.sql.BasicBinder`) had to be updated to TRACE in the `log4j.xml` file to enable logging of SQLs as well as bind parameters:


```

<!-- Use "TRACE" to see the debug output from UDA and activate tracing in stored
procedures.
    Check also server.properties property "db.default.debug.show_sql" -->
<category name="UDA" additivity="false">
  <priority value="TRACE"/>
  <appender-ref ref="UdaAppender"/>
</category>

<!-- Use "TRACE" to see the debug output from Hibernate
    org.hibernate.cache           Log all second-level cache activity
    org.hibernate.hql.ast.AST     Log HQL and SQL ASTs during query parsing
    org.hibernate.jdbc            Log all JDBC resource acquisition
    org.hibernate.pretty          Log the state of all entities (max 20 entities)
associated with the session at flush time
    org.hibernate.SQL             Log all SQL DML statements as they are executed
(include org.hibernate.type to see binds and values)
    org.hibernate.secure          Log all JAAS authorization requests
    org.hibernate.tool.hbm2ddl    Log all SQL DDL statements as they are executed
    org.hibernate.transaction    Log transaction related activity
    org.hibernate.type            Log all JDBC parameters
    org.hibernate                 Log everything -->
<category name="org.hibernate.SQL" additivity="false">
  <level value="TRACE" />
  <appender-ref ref="UdaAppender"/>
</category>
<category name="org.hibernate.type.descriptor.sql.BasicBinder" additivity="false">
  <level value="TRACE" />
  <appender-ref ref="UdaAppender"/>
</category>

```

Changes in this file do not require a server restart, and after the above query is executed again, you will need to find the query in the logs.

Here's the query after we've replaced the necessary bind parameters:

```

set linesize 1000

explain plan for
select articlerev0_.ID, articlerev0_.Identifier
  from
    "ArticleRevision" partition(P_PRODUCT) articlerev0_
  left outer join
    "ArticleDetail" partition(P_PRODUCT) articledet1_
    on articlerev0_.ID=articledet1_.ArticleRevisionID
    and articledet1_.DeletionTimestamp = TO_TIMESTAMP('9999-12-31','yyy
y-mm-dd')
  left outer join
    "ArticleLang" partition(P_PRODUCT) articlelan2_
    on articlerev0_.ID=articlelan2_.ArticleRevisionID
    and (
      articlelan2_.LanguageID=9
      and articlelan2_.Res_LK_Text100_01='LT'
      and articlelan2_.Res_LK_Int_01=0
      and articlelan2_.ChannelID=1
      and articlelan2_.EntityID=21007
    )
    and articlelan2_.DeletionTimestamp = TO_TIMESTAMP('9999-12-31','yyy
y-mm-dd')
  left outer join
    "ArticleAttribute" partition(P_PRODUCT) articleatt3_
    on articlerev0_.ID=articleatt3_.ArticleRevisionID
    and (
      articleatt3_.NameInKeyLanguage='Heel Type'
    )
    and articleatt3_.DeletionTimestamp = TO_TIMESTAMP('9999-12-31','yyy
y-mm-dd')
  left outer join
    "ArticleAttributeValue" partition(P_PRODUCT) articleatt4_
    on articleatt3_.ID=articleatt4_.ArticleAttributeID
    and (
      articleatt4_.Identifier='ALL'
      and articleatt4_.LanguageID=9
    )
    and articleatt4_.DeletionTimestamp = TO_TIMESTAMP('9999-12-31','yyy
y-mm-dd')
  where
    articlerev0_.RevisionID = 1
    and articlerev0_.DeletionTimestamp = TO_TIMESTAMP('9999-12-31','yyyy-
mm-dd')

    and articledet1_.Res_Bit_02=1
    and articlelan2_.Res_Text250_05='41'
    and (
      articleatt4_.Value is not null
    )
    and articlerev0_.EntityID=1100;

```

Explaining the plan query for this resulted in the following. Notice the difference between the estimated row count and the actual row count? This results in Oracle choosing an incorrect plan with an incorrect join order and many nested loops.



Also notice that an INDEX RANGE SCAN was performed on index XAK2_ArticleAttribute in table ArticleAttribute, which resulted in an incorrect estimated row count.

To investigate further, a 10053 Optimizer trace is needed. In the same session from where you issued the query and explain plan, do the following. Pay attention to the commented lines below. You'll need to turn on tracing, execute the query, then turn off tracing. Then you'll need to view the trace files, preferably with a utility:

```
-- activate generating trace files
ALTER SESSION SET EVENTS='10053 trace name context forever, level 1';
-- deactivate generating trace files
ALTER SESSION SET EVENTS '10053 trace name context off';
-- get filenames for trace session
select u_dump.value || '/' || instance.value || '_ora_' || v$process.spid
|| nvl2(v$process.traceid, '_' || v$process.traceid, null) || '.trc'"Trace File"
from V$PARAMETER u_dump
cross join V$PARAMETER instance
cross join V$PROCESS
join V$SESSION on v$process.addr = V$SESSION.paddr
where u_dump.name = 'user_dump_dest'
and instance.name = 'instance_name'
and V$SESSION.audsid=sys_context('userenv','sessionid');
```

Navigate to the Explain Plan section of the trace file to confirm you are looking at the proper query in the trace file:

```
----- Explain Plan Dump -----
```

```
----- Plan Table -----
```

```
=====
```

```
Plan Table
```

```
=====
```

+-----+-----+						
Id	Operation				Name	Rows
Bytes	Cost	Time	Pstart	Pstop		
+-----+-----+						
0	SELECT STATEMENT					
10						
1	NESTED LOOPS					1
116	10	00:00:01				
2	NESTED LOOPS					1
116	10	00:00:01				
3	NESTED LOOPS					1
84	7	00:00:01				
4	PARTITION REFERENCE SINGLE					1
55	5	00:00:01	KEY(AP)	KEY(AP)		
5	TABLE ACCESS BY LOCAL INDEX ROWID				ArticleAttribute	1
55	5	00:00:01	2	2		
6	INDEX RANGE SCAN				XAK2_ArticleAttribute	1
4	00:00:01	2	2			
7	TABLE ACCESS BY GLOBAL INDEX ROWID				ArticleRevision	1
29	2	00:00:01	2	2		
8	INDEX UNIQUE SCAN				PK_ArticleRevision	1
1	00:00:01					
9	PARTITION REFERENCE SINGLE					1
2	00:00:01	KEY(AP)	KEY(AP)			
10	INDEX UNIQUE SCAN				XAK1_ArticleAttributeValue	1
2	00:00:01	KEY(AP)	KEY(AP)			
11	TABLE ACCESS BY LOCAL INDEX ROWID				ArticleAttributeValue	1
32	3	00:00:01	1	1		
+-----+-----+						

In my trace file the explain plan dump for the query starts at line 6613.

Using a trace file viewer is recommended - general steps to find the issue are:

1. Find the proper 'optimizing query block' (OPTIMIZING QB) section for your sql (If using v10053.exe look for the OPTIMIZING QB section on the left hand side).
2. Look for the Access path analysis for the table containing the wrong statistic (ArticleAttribute in our case).
3. Look for the column containing the index giving the wrong estimated row count. In our case, the XAK2_ArticleAttribute index is on the NameInKeyLanguage column.

At this point you'll want to familiarize yourself with how Oracle uses histograms to manage statistics.



If you look closely at section 3, you'll notice that the NDV (Number of Distinct Values) is 383, but the #Bkts (Buckets) is 255 from the HtBal (Height Balanced) histogram. The NameInKeyLanguage value of 'Heel Height' was unfortunately not in one of the buckets, which causes the "Using prorated density: 0.000000 of col #54 as selectvity of out-of-range/non-existent value pred" message you see. This causes the optimizer to incorrectly guess that this value doesn't exist in this table and that joining to this table would also return no rows for the query (prorated density of 0).

To fix this, we need to delete the histogram. Here's a script which deletes all histograms for all varchars with > 254 distinct values, as the number of buckets isn't sufficient to cover all values so a height balanced histogram isn't desired for this scenario (in general, height balanced histograms are almost useless for varchar2 columns).

Script which deletes all height balanced histograms for nvarchar2 columns with more than 254 distinct values

```
-- script which deletes all height balanced histograms for nvarchar2 columns with
more than 254 distinct values
BEGIN
  FOR COL_ITEM IN (
    SELECT stats.OWNER, stats.TABLE_NAME, stats.COLUMN_NAME
    FROM DBA_TAB_COL_STATISTICS stats
    INNER JOIN DBA_TAB_COLUMNS cols
    ON stats.TABLE_NAME = cols.TABLE_NAME
    AND stats.COLUMN_NAME = cols.COLUMN_NAME
    WHERE cols.DATA_TYPE = 'NVARCHAR2'
    AND stats.HISTOGRAM <> 'NONE'
    AND stats.NUM_DISTINCT > 254
    AND stats.OWNER = cols.OWNER
    AND stats.OWNER IN ('HBC_HPM_MAIN', 'HBC_HPM_MASTER', 'HBC_HPM_SUPP
LIER')
    ORDER BY stats.OWNER, stats.TABLE_NAME, stats.COLUMN_NAME
  )
  LOOP
    DBMS_STATS.DELETE_COLUMN_STATS (OWNNAME=>COL_ITEM.OWNER, TABNAME=>' ' ||
COL_ITEM.TABLE_NAME || ' ', COLNAME=>' ' || COL_ITEM.COLUMN_NAME || ' ',
COL_STAT_TYPE=>'HISTOGRAM');
  END LOOP;
END;
/
```

Finally, to prevent Oracle's automatic statistics job from regenerating the histogram you need to call the set_table_prefs function.

Pass size 1 as the METHOD_OPT to prevent histogram creation:

```
--This does not work as column and table name are lowercase and must be passed with
double quote.
--TODO is to fix this call to work with mixed case table and column names
BEGIN
  DBMS_STATS.set_table_prefs('HBC_HPM_MASTER', '"ArticleAttribute"', 'METHOD_OPT',
    'for all columns size skewonly for columns size 1 "NameInKeyLanguage"');
END;
/
```

After fixing this and a similar issue in the Res_Text250_05 column for table ArticleLang, we get the following plan. While the actual and estimated row counts are off, the join order is now more efficient (Good: ArticleLang, ArticleDetail, ArticleAttribute, ArticleAttributeValue vs. Bad: ArticleAttribute, ArticleDetail, ArticleLang, ArticleAttributeValue previously). For outer joins in general, the most limiting tables should be joined first, which for PIM means ArticleAttribute and ArticleAttributeValue should usually be joined last as they are larger than the Article table by some factor.



1.5.1 Fixing frequency histogram issues

The second issue look for the Res_Text250_05 column for table ArticleLang:

```
Column (#13):
NewDensity:0.000053, OldDensity:0.000000 BktCnt:9392, PopBktCnt:9375, PopValCnt:114,
NDV:208
Column (#13):
NewDensity:0.000122, OldDensity:0.000000 BktCnt:4113, PopBktCnt:4091, PopValCnt:103,
NDV:208
Column (#13): Res_Text250_05( Part#: 1
AvgLen: 4 NDV: 208 Nulls: 10118996 Density: 0.000122
Histogram: Freq #Bkts: 125 UncompBkts: 4113 EndPtVals: 125
Column (#13): Res_Text250_05(
AvgLen: 4 NDV: 208 Nulls: 10118996 Density: 0.000122
Histogram: Freq #Bkts: 125 UncompBkts: 4113 EndPtVals: 125
...
...
Using density: 0.000122 of col #13 as selectivity of unpopular value pred
```

Issue: The value selected did not fall into any of the histogram buckets.

Solution: Add more buckets such that all values in the column can be assigned a bucket by forcing the optimizer to analyze 100 percent of the rows. As the number of distinct values in this column is less than 254 (Oracle max # of buckets) this is a suitable solution:

```

-- script which adds more buckets for frequency histograms
DECLARE
    BUCKETS_BUFFER NUMBER := 0;
BEGIN
    FOR COL_ITEM IN (
        SELECT stats.OWNER, stats.TABLE_NAME,
            LISTAGG('FOR COLUMNS "' || stats.COLUMN_NAME || '" SIZE ' || CASE
                WHEN stats.NUM_DISTINCT + BUCKETS_BUFFER > 254 THEN 254 ELSE stats.NUM_DISTINCT +
                BUCKETS_BUFFER END, ', ') WITHIN GROUP (ORDER BY stats.COLUMN_NAME) AS METHOD_OPT
        FROM DBA_TAB_COL_STATISTICS stats
            INNER JOIN DBA_TAB_COLUMNS cols
                ON stats.TABLE_NAME = cols.TABLE_NAME
                AND stats.COLUMN_NAME = cols.COLUMN_NAME
        WHERE cols.DATA_TYPE = 'NVARCHAR2'
        AND stats.HISTOGRAM = 'FREQUENCY'
        AND stats.NUM_DISTINCT > stats.NUM_BUCKETS
        AND stats.OWNER = cols.OWNER
        AND stats.OWNER IN ('HBC_HPM_MAIN', 'HBC_HPM_MASTER', 'HBC_HPM_SUPP
LIER')


        GROUP BY stats.OWNER, stats.TABLE_NAME
        ORDER BY stats.OWNER, stats.TABLE_NAME
    )
    LOOP
        DBMS_STATS.GATHER_TABLE_STATUS
        (
            OWNNAME => COL_ITEM.OWNER,
            TABNAME => '"' || COL_ITEM.TABLE_NAME || '"',
            ESTIMATE_PERCENT => 100,
            METHOD_OPT => COL_ITEM.METHOD_OPT,
            CASCADE => FALSE
        );
    END LOOP;
END;
/

```

1.6 Oracle Import/Export for MAIN, MASTER and SUPPLIER Schemas

A "how to" guideline for exporting and importing Product Manager Schemas from and to an Oracle 11G. Informatica Software Development or Support might ask you to provide the data of a certain installation in order to reproduce bugs or performance issues. This guide shall help support, partners and consulting to provide this.

- [Preparation](#)(see page 40)
 - [Directory Object](#)(see page 40)
- [Export](#)(see page 41)
 - [User Roles](#)(see page 41)
 - [Execute export datapump](#)(see page 41)
- [Import](#)(see page 41)
 - [User Roles](#)(see page 41)
 - [Tablespaces](#)(see page 42)
 - [Execute Import Datapump](#)(see page 43)

 The terms backup/restore which are usually used in the SQL Server world map to export/import in Oracle. The actual backup/restore functionality of Oracle is bound to a specific oracle instance (more or less) and is only meant for error/data loss recovery, but not to transfer the data to a different database instance.

1.6.1 Preparation

The whole export/import process can not be performed by the sys user at all. Therefore we suggest to create a new user for administrative purposes. In the following sections we will refer to this user as `hlradmin`. The effective permissions a user needs to have differ between the export and import process, thus they are described in the corresponding sections.

Create the hlradmin user

1	<pre>CREATE USER "HLRADMIN" IDENTIFIED BY "heiler" PROFILE "DEFAULT" ACCOUNT UNLOCK;</pre>
---	---

1.6.1.1 Directory Object

The directory object in oracle maps to a physical directory path. This directory is used to store the dumpfile(s). The directory object is needed for the export as well as the import.

Create the directory object

1	CREATE DIRECTORY "HLR_DIR" AS 'c:\heiler\oracle\impexp'
---	---

1.6.2 Export

1.6.2.1 User Roles

The hlradmin user needs the following privileges to export data

- resource
- connect
- select permissions on all objects for the MAIN, MASTER and SUPPLIER users

Grant needed privileges to hlradmin user

1	GRANT CONNECT TO "HLRADMIN";
2	GRANT RESOURCE TO "HLRADMIN";
3	GRANT EXP_FULL_DATABASE TO "HLRADMIN";

1.6.2.2 Execute export datapump

1.6.3 Import

1.6.3.1 User Roles

The hlradmin user needs the following privileges to perform the import:

- imp_full_database
- create user
- create any table
- create any index
- create any sequence
- create any procedure
- create tablespace

Grant needed privileges to hlradmin user

```

1  GRANT IMP_FULL_DATABASE TO "HLRADMIN";
2  GRANT CREATE USER TO "HLRADMIN";
3  GRANT CREATE ANY TABLE TO "HLRADMIN";
4  GRANT CREATE TABLE WITH ADMIN OPTION TO "HLRADMIN";
5  GRANT CREATE ANY INDEX TO "HLRADMIN";
6  GRANT CREATE INDEX WITH ADMIN OPTION TO "HLRADMIN";
7  GRANT CREATE ANY SEQUENCE TO "HLRADMIN";
8  GRANT CREATE SEQUENCE WITH ADMIN OPTION TO "HLRADMIN";
9  GRANT CREATE ANY PROCEDURE TO "HLRADMIN";
10 GRANT CREATE PROCEDURE WITH ADMIN OPTION TO "HLRADMIN";
11 GRANT CREATE TABLESPACE TO "HLRADMIN";

```

1.6.3.2 Tablespaces

To import a dump file which contains the tablespaces of the MAIN, MASTER and SUPPLIER schemas, you need to have corresponding tablespaces in the target database. You can execute the following pl/sql snippet for each needed schema in order to create the tablespaces.

Create empty tablespaces

```

1  -- Temporary tablespace
2  CREATE TEMPORARY TABLESPACE "<SCHEMA_NAME>_TEMP"
3      TEMPFILE 'c:/heiler/oracle/impexp/<SCHEMA_NAME>_Temp.dbf'
4      SIZE 1024M REUSE
5      AUTOEXTEND ON NEXT 128M
6      MAXSIZE UNLIMITED
7      EXTENT MANAGEMENT LOCAL
8      UNIFORM SIZE 1M;
9  -- Data tablespace
10 CREATE TABLESPACE "<SCHEMA_NAME>_DATA"
11     DATAFILE 'c:/heiler/oracle/impexp/<SCHEMA_NAME>_Data.dbf'
12     SIZE 1024M REUSE
13     AUTOEXTEND ON NEXT 128M
14     MAXSIZE UNLIMITED
15     EXTENT MANAGEMENT LOCAL
16     SEGMENT SPACE MANAGEMENT AUTO;
17 -- Separate index tablespace
18 CREATE TABLESPACE "<SCHEMA_NAME>_INDEX"
19     DATAFILE 'c:/heiler/oracle/impexp/<SCHEMA_NAME>_Index.dbf'
20     SIZE 1024M REUSE
21     AUTOEXTEND ON NEXT 128M
22     MAXSIZE UNLIMITED
23     EXTENT MANAGEMENT LOCAL
24     SEGMENT SPACE MANAGEMENT AUTO;

```

1.6.3.3 Execute Import Datapump

Parameter file

```

1  # the directory object
2  DIRECTORY=HLR_DIR
3  # the name of the ora dump file, located in the physical file directory
   mapped to DIRECTORY
4  DUMPFILE=MYEXPORT01.DMP
5  # the schemas to be imported
6  # attention: you have to import all PIM - Server schemas (there're
   dependencies between the schema objects)
7  SCHEMAS=CUST_MAIN,CUST_MASTER,CUST_SUPPLIER
8  # specify the schema name mappings, necessary if schema names to be
   imported are different from local schema names
9  # schemas are created automatically if they don't exist
10 # pattern:
11 # schema name in dump file:local schema name
   REMAP_SCHEMA=CUST_MAIN:HPM_MAIN,CUST_MASTER:HPM_MASTER,CUST_SUPPLIER:HPM_S
   UPPLIER
12 # specify the tablespace name mappings, necessary if tablespace names to
   be imported are different from local tablespace names
13 # tablespace are NOT created automatically, you have to create them
   before
14 # pattern:
15 # tablespace name in dump file:local tablespace name (must exist)
16 REMAP_TABLESPACE=
   CUST_MAIN_INDEX:HPM_MAIN_INDEX,CUST_MAIN_DATA:HPM_MAIN_DATA,CUST_MAIN_TEMP
   :HPM_MAIN_TEMP,CUST_MASTER_INDEX:HPM_MASTER_INDEX,CUST_MASTER_DATA:HPM_MAS
   TER_DATA,CUST_MASTER_TEMP:HPM_MASTER_TEMP,
   CUST_SUPPLIER_INDEX:HPM_SUPPLIER_INDEX,CUST_SUPPLIER_DATA:HPM_SUPPLIER_DAT
   A,CUST_SUPPLIER_TEMP:HPM_SUPPLIER_TEMP
17 # use this option to only create the sql file
18 # no import will be executed
19 # SQLFILE=SQL_IMPORT.SQL

```

execute import cmd

```

1  impdp USERID=HLRADMIN/heiler parfile="<path to parameter file>" | pause


```

i If you get errors regarding insufficient access rights of user HLRADMIN, try to execute with SYSTEM-User:

```

impdp USERID=SYSTEM/heiler parfile="<path to parameter file>" | pause

```

 While importing, you might get several Oracle errors regarding the users .._HPM_BULKLOAD and .._HPM_REPSTORE:

ORA-39083: Objekttyp OBJECT_GRANT konnte nicht erstellt werden, Fehler:

ORA-01917: Benutzer oder Funktion KRAMP_HPM_REPSTORE ist nicht vorhanden

Fehlerhafte SQL ist:

```
GRANT FLASHBACK ON "KRAMP_HPM_MAIN"."User" TO "KRAMP_HPM_REPSTORE"
```

ORA-39083: Objekttyp OBJECT_GRANT konnte nicht erstellt werden, Fehler:

ORA-01917: Benutzer oder Funktion KRAMP_HPM_BULKLOAD ist nicht vorhanden

Fehlerhafte SQL ist:

```
GRANT ALTER ON "KRAMP_HPM_MAIN"."UnitSystemUnitMapLang" TO "KRAMP_HPM_BULKLOAD"
```

They do not have any impact on the import result, so you can ignore them.

2 Server Operation

Combines maintenance, tuning and troubleshooting for the Product 360 application server.

- [Maintenance Tasks](#)(see page 45)
- [Monitoring](#)(see page 45)
- [Tuning](#)(see page 46)
 - [Memory Settings](#)(see page 46)
 - [Garbage Collection](#)(see page 46)
 - [GC activity log](#)(see page 46)
 - [Analyse GC activity log & optimize](#)(see page 47)
 - [Communication Framework](#)(see page 47)
 - [Mass-Data Parallel Processing](#)(see page 47)
 - [Media Asset Parallel Management](#)(see page 47)
 - [CreateZipManagement](#)(see page 47)
 - [MediaAssetProvider special MBean](#)(see page 49)
 - [MediaAssetProviderParallelManager](#)(see page 49)
 - [HeilerClassic special MBeans](#)(see page 50)
 - [CreateThumbnailManager](#)(see page 50)
 - [HeilerClassicParallelManager](#)(see page 50)
 - [Database Connection Pool](#)(see page 50)
 - [Logging](#)(see page 51)
 - [Import](#)(see page 52)
 - [Preferences regarding CPU usage](#)(see page 52)
 - [Preferences regarding memory usage](#)(see page 53)
- [Trouble Shooting](#)(see page 54)
 - [Anti-Virus Software affecting performance](#)(see page 54)
 - [Media Asset Provider \(Classic Provider\)](#)(see page 54)
 - [No preview are visible](#)(see page 54)
 - [Obtain a Memory Dump \(Heap Dump\) from the Java Virtual Machine](#)(see page 55)
 - [Enhanced Logging for Assortment Content Calculation](#)(see page 55)
 - [Server and client in different time zones](#)(see page 56)
 - [Cluster](#)(see page 56)

2.1 Maintenance Tasks

- Backup Strategy
- Background job should be configured to be automatically deleted after a certain amount of time. See [Server Job Maintenance](#)(see page 70) for details.

2.2 Monitoring

Monitoring tools which can communicate using the SNMP protocol can be used to monitor the Product 360 application servers. Please see [SNMP Monitoring](#)(see page 57) for configuration details on the SNMP functionality.


Additionally micrometer metrics are available. See page [Monitoring with Micrometer](#)(see page 86) for more information.

2.3 Tuning

Here is the list of settings and configuration options that may affect system performance. These settings depend on system workloads, hardware and environment.


2.3.1 Memory Settings

As expected, the memory settings are crucial for the application server of the PIM - Server. You can configure the available memory in the `<installation root>/_environment.conf` file. Please note that you can configure the maximum amount of memory (`MEM_MAX`).

 The minimum amount of available heap memory will always be set to the same value as the `MEM_MAX`, since for an application server it makes no sense otherwise.

2.3.2 Garbage Collection

Garbage collection configuration is a very complicated thing to do. Informatica PIM Software engineers spend a lot of time on site with customers to find the right garbage collection algorithm and settings. All needed parameters have already been adjusted by us, the only configuration which is needed for every customer installation is the number of garbage collection threads which can be used. PIM - Server uses a parallel garbage collection algorithm, therefore you can decide how many threads the garbage collection can use. The `GC_THREADS` setting in the `<installation root>/_environment.conf` file should be set to about 75% of the available CPU cores. For example, in case you have 8 Cores (including hyper-threading cores), you would set this to about 6 threads.

 In case you need to further adjust the garbage collection settings you can do so in the `<installation root>/service/wrapper.conf` file in the `GC_OPTIONS` section. Please note that the `console_debug.cmd` launch file will use different garbage collection algorithms since the debug mode is not able to handle the ones we use for production.

2.3.2.1 GC activity log

In order to analyze the GC events in detail you can use the GC activity log for each Product 360 Server located in the `<installation root>/logs` folder. GC activity logs are enabled per default for all Product 360 Servers.

Each Server start produces a new `gc_<timestamp>.log.*` file (whereas `<timestamp>` is the starting time of the Product 360 Server) in the server's log folder.

2.3.2.2 Analyse GC activity log & optimize

The GC activity log can be analyzed e.g. online via <http://gceasy.io/>. Common issues and optimization possibilities are:

- The "GC pause duration time" specifies the length of the time intervals, where the garbage collector runs exclusively and therefor pauses all other application threads. If the **"GC pause duration time" is above 10 seconds**, the Product 360 Server application runs into trouble, because there are continuously running heartbeats (e.g. for the Hazelcast Framework, which have a timeout of 10 seconds. This can result in operation timeouts in the log file (e.g. HazelcastOperationTimeout)
 - **For optimization: consider increasing heap memory** (MEM_MAX in <installation root>/_environment.conf file)
- The "object creation rate" (for <http://gceasy.io/>: see table "Object stats", row "Avg creation rate") specifies the heap memory size, which is allocated per second by new created java objects. If there are problems with operation timeout exceptions in the log due to long "GC pause duration time"(see the explained issue above) AND the **"object creation rate" is also very high** (somewhere above 1 GB/s), then the garbage collector is maybe running too late and then is not able to free memory in less than 10 seconds. By default the garbage collector G1 begins to work, when 45% of the heap memory is in use. This is specified by the parameter "InitiatingHeapOccupancyPercent", which has a default value of 45.
 - **For optimization: consider decreasing the "InitiatingHeapOccupancyPercent"**. Therefor add -XX:InitiatingHeapOccupancyPercent=30 (or even 20 or 10) to GC OPTIONS section in <installation root>/service/wrapper.conf

Find general information about the used Garbage collection algorithm G1 here: <https://www.oracle.com/technetwork/tutorials/tutorials-1876574.html>

2.3.3 Communication Framework

Number of worker threads in server network settings

2.3.4 Mass-Data Parallel Processing

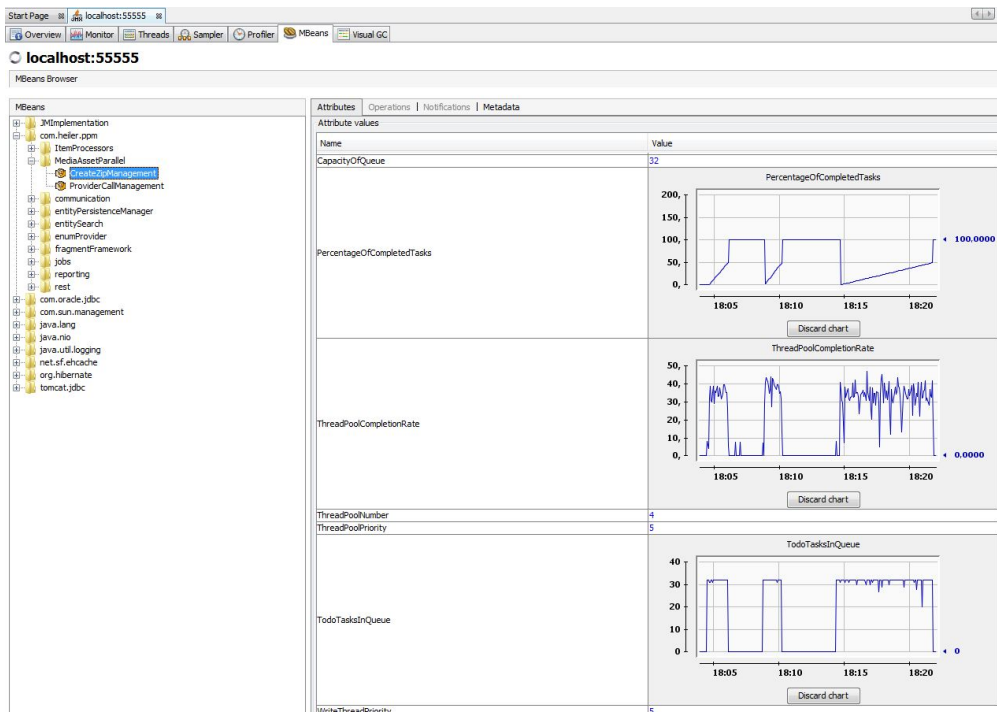
Item Processor Threadpool

2.3.5 Media Asset Parallel Management


There are several registered MBeans under the category **com.heiler.ppm** → **MediaAssetParallel** with the JMX interface which are applied to observe and improve the performance in run time.

2.3.5.1 CreateZipManagement

It provides the management of the thread pool and other relevant parameters(e.g. the size of the BlockingQueue to avoid OutOfMemoryError) by zip creation phase.



The following screenshot indicates the corresponding threads in run time.



2.3.5.2 MediaAssetProvider special MBean

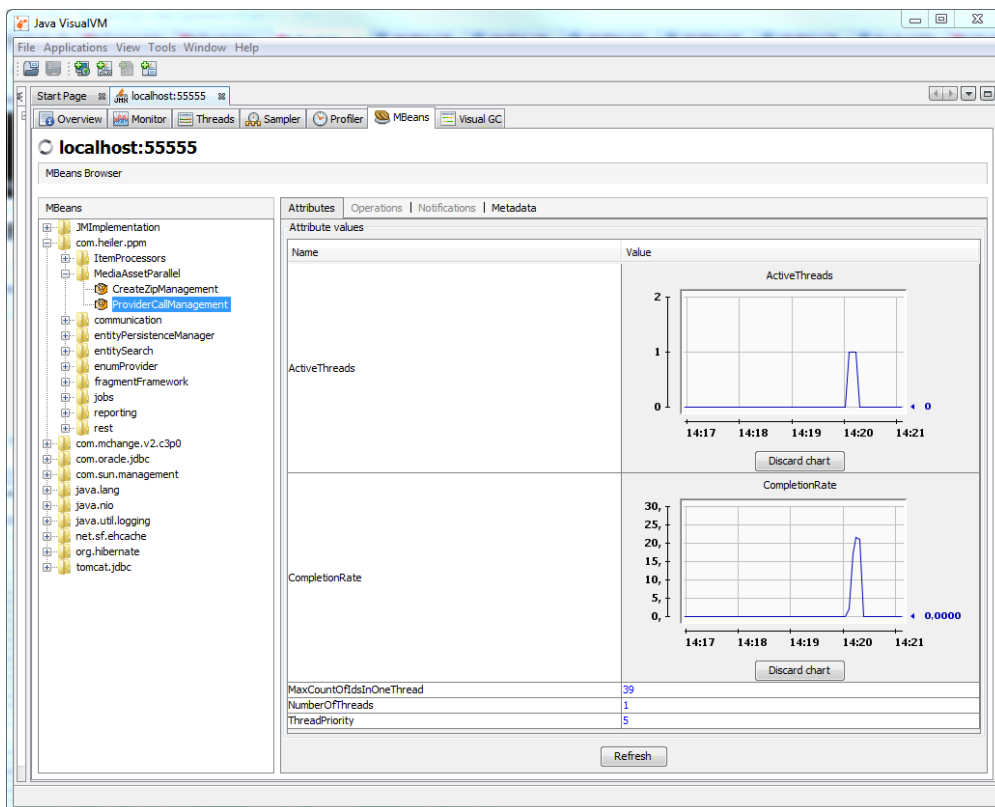
MediaAssetProviderParallelManager

i It was named as "ThreadPoolManagement", and changed as "ProviderCallManagement" and then as "MediaAssetProviderParallelManager" since 7.1.01.00.

It provides the management of the multi-threads call for the corresponding MediaAssetProvider methods.

A special thread pool is applied to improve the performance for some methods of the default IMediaAssetProvider, such methods should have any parameter composed of a long id list (e.g. getUNCpaths(String[] identifiers, ...)). Currently only the PIM - Media Manager 'HMM' MediaAssetProvider supports it. This thread pool settings can be observed and adjusted with the JMX interface during runtime. It is meanful for some scenarios like export with mass data.

The following screenshot indicates the thread pool utilization. The user can adjust the values of the attributes "MaxCountOfIdsInOneThread" and "NumberOfThreads" to improve the performance, if the "ActiveThreads" reaches always the maximum value and the "CompletionRate" is but too low.



2.3.5.3 HeilerClassic special MBeans

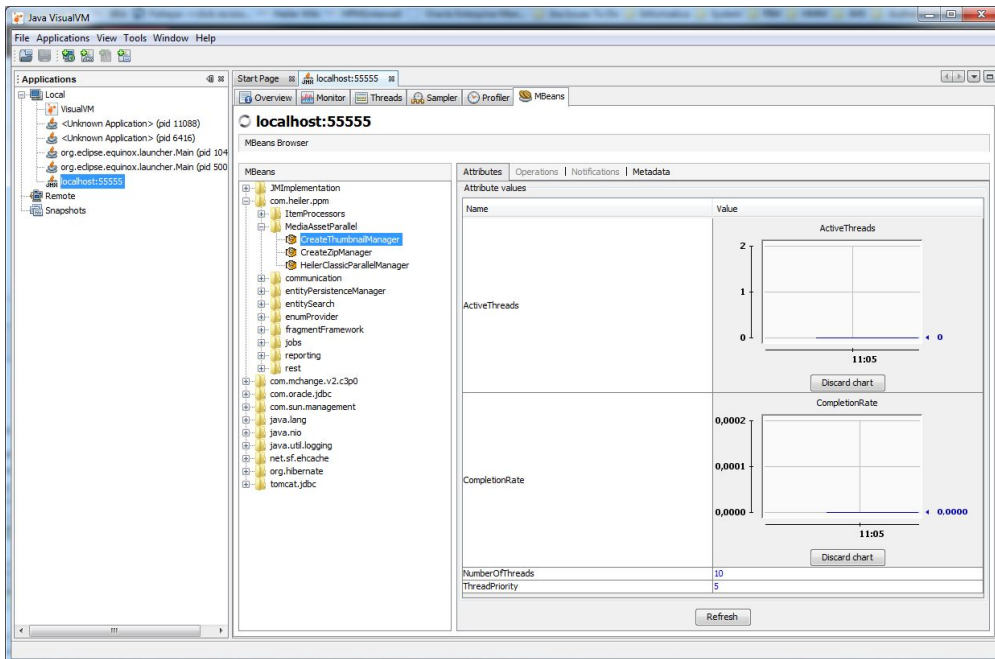
i since 7.1.01.00 HeilerClassic registers two MBeans which can be also observed and adjusted with JMX interface.

CreateThumbnailManager

It provides the management of thread pool which schedules all graphicsmagick convert jobs to create thumbnails. The default NumberOfThread is 10.

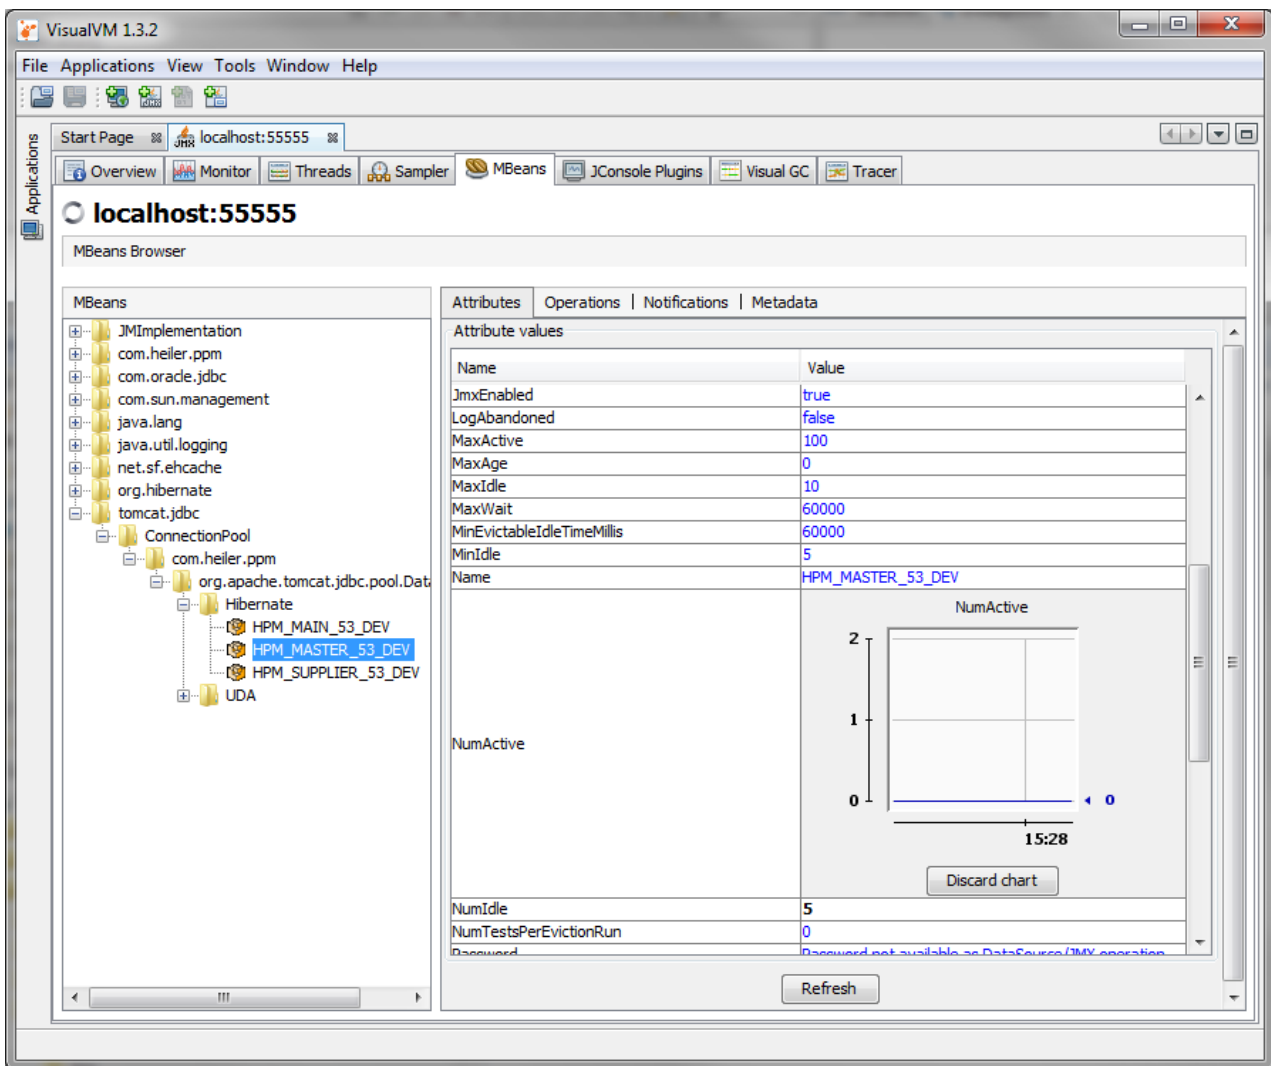
HeilerClassicParallelManager

It provides the management of thread pool which schedules tasks which can be performed parallelly for the corresponding MediaAssetProvider methods(e.g. GetSubCategoriesTask).



2.3.6 Database Connection Pool

The database connection pool settings can be adjusted in the corresponding configuration files. More detailed information on the configuration of the connection pool can be found here [DB Connection pool](#)(see page 65). During runtime the connection pool utilization can be observed with the JMX interface (see Screenshot). Most important here is the maximum used attribute which indicates if the maximum number of connections has already been used. You might need to increase the maximum connections in case the limit is reached often, or continuously.



This screenshot shows the NumActive attribute of the ConnectionPool for the HPM_MASTER connection pool. Currently there are zero connections in use. When this number reaches the MaxActive limit, the application load might be too high for the number of connections, adjusting the connection pool is necessary then. Additionally to NumActive and MaxActive you might want to check the WaitCount, which indicates how long a thread needs to wait until it acquires a connection from the pool. In case there are enough connections left in the pool, this number should be about zero.

⚠ The PIM - Server has two connection pool configurations which are of interest here, the Hibernate and the UDA one. The Hibernate pool is used for detail model persistence, therefore it's utilized in detail model modifications, import and merge. The UDA pool is used mostly for read-only mass-data retrieval like list model based functionality such as table views and export.

2.3.7 Logging

PIM - Server has a central configuration file for all logging activities in the server, as well as in the client. By default the logging trace levels are set to INFO, thus logging messages with severity INFO, WARNING and ERROR will be logged in the console log, errors additionally in the error.log file. Please make sure that TRACE

or DEBUG severity levels are disabled in production environments, since they dramatically slow down the performance.

Logging is configured in the <installation root>/configuration/hpm/log4j.xml file.

2.3.8 Import

2.3.8.1 Preferences regarding CPU usage

Especially in Single Server scenarios in which the Import Jobs are executed throughout the day, it might be necessary to limit the CPU resource consumption of a single import process.

For this there are several preferences available which can be set in the plugin_customization.ini file.

```
# Specifies the maximum number of threads used for importing objects.
```

```
# Default is twice the number of cores.
```

```
# Allowed value is an integer greater than 0.
```

```
# com.heiler.ppm.importer5.core/importer.maxThreads=2
```

```
# Factor for the number of parallel threads.
```

```
# The number of parallel threads is computed by multiplying the number of cores and this factor.
```

```
# The result is rounded.
```

```
# Value is a double; default is 2.0
```

```
# This value is only considered if importer.maxThreads is not set.
```

```
com.heiler.ppm.importer5.core/importer.parallelThreadsFactor = 2.0
```

```
# Allow to limit number of the import jobs running in parallel
```

```
# Only the actual database import (phase 2) is limited.
```

```
# -1 means unlimited (default)
```

```
com.heiler.ppm.importer5.core/importer.maxNumberOfDatabaseImportsAtSameTime = -1
```

Additionally to these global settings, a user can activate the so called "debug mode" for the Import by pressing

Shift + F12 when **inside of the Import Perspective** in the Desktop Client. With activated debug mode, additional settings are available when scheduling the import.

Import

Schedule import
Specify when the import is to be executed.

As soon as possible ▾

Time ▾

Structure: Heiler Standard ▾

Import mode: Import or update all objects ▾

Error mode: Tolerant (object is imported as far as possible) ▾

☐ Perform test run before import, cancel import on errors

Test Settings

☐ Generate gold standard file

☐ Test against gold standard file

Number of threads:

☐ Use always bulkload

☐ Use never bulkload

Bulkload size:

☐ Clear catalog

Raw performance test (number of objects):

Storage mode for intermediate memory (0-2):

OK Cancel


⚠ Please do not change any of those additional settings except the Number of Threads!

2.3.8.2 Preferences regarding memory usage

When importing CSV or Excel files, per default all rows are read into heap memory in order to presort them. This speeds up processing if the data for one item is stored in rows that are spread. So when importing large CSV or Excel data files, consider to increase the job servers' heap size accordingly.

If presorting is not needed (because the data files are already presorted) or desired, it can be disabled by adding the following preference in the server's `plugin_customization.ini`:

`com.heiler.ppm.mapper.core/presortRowsDefault = false`

 Importing large data files might become slower, when presorting is deactivated (see preference above) and the import data file is not already presorted.

2.4

Trouble Shooting

- Log files
- What's needed for a support call

2.4.1 Anti-Virus Software affecting performance

In the past we had some customer performance issues which were bound to the anti-virus scanning on the application server. Although we can not say that this is always the case for all deployments or all anti-virus products, we do recommend to exclude the PIM relevant directories from the scanning if applicable by the customers local security policy. If scanning can't be avoided in your situation, you should consider to reduce the lifetime of import or export jobs. The shorter they are, the less files will be in the import/export directories - which might increase scanning performance again.

2.4.2 Media Asset Provider (Classic Provider)

2.4.2.1 No preview are visible

Ensure following things:

1. Is the file extension unsupported?(E.g. pdf, ps, eps files because the ghostscript is not supported in PIM any more)
2. Is the graphicsMagick(official verified version GraphicsMagick-1.3.14-Q16) installed?
3. Check that the `blacklistExtensions` parameter of the `C:\heiler\server\configuration\HPM\plugin_customization.ini` file does not contain the file extension
4. Has the user who has started the application server write access to the media folder defined in the "Media Asset Server (MAS) Settings" section of the `server.properties` file?
Is it possible to open the image files from another computer by the explorer and another program?
5. Important: An unsuccessful rendering of a thumbnail leads to an entry in the database table `MediaAssetFile` of the "[db_prefix]_MAIN_[db_suffix]" database.
Therefore, after a "subsequent" adjustment of these points this table must be modified. Set the bit-field "SupportThumbnail" to NULL.

Check 3rd party tools:

The technical background is that PIM - Server runs on a 64-bit machine as a 64-bit Java process. PIM - Server starts `gm.exe` (Graphicsmagick) in a several process space (Graphicsmagick is a 32-bit process).

2.4.3 Obtain a Memory Dump (Heap Dump) from the Java Virtual Machine

There are several ways to obtain a heap dump during runtime.

- Heap dump on out of memory
By default all Product Manager launch configurations (development and runtime) contain the Java parameter `-XX:+HeapDumpOnOutOfMemoryError` which performs a heap dump when an out of memory error occurs.
- Heap dump on Ctrl+Break
Beginning with Java Version 1.5.0_16 the Java VM supports the `-XX:+HeapDumpOnCtrlBreak` parameter which will create a heap dump when you press Ctrl + Break in the Java console. For this the PIM - Server must be started with the `debug_console.cmd` file.
- Heap dump via jcmd utility
This console utility is included in JDK. Use `"jcmd <processId> GC.heap_dump <filename>"` for creating a heap dump, e.g. `"jcmd 14904 GC.heap_dump D:\heapdump.hprof"`
For details see <https://docs.oracle.com/javase/8/docs/technotes/guides/troubleshoot/tooldescr006.html>.

We suggest to use the Eclipse Memory Analyzer to analyze the Heap Dump, especially in case you have large heap sizes. It's the only tool we know so far which is able to open large heap dumps at all (you will need a lot of memory during the indexing of the heap dump in the memory analyzer - but this must not be physical memory 😊)

2.4.4 Enhanced Logging for Assortment Content Calculation

In case there are issues in terms of an assortment's content and there are items included/excluded which do not fit to the defined assortment rules, please add the following lines to the server's `log4j.xml` and execute an assortment update:

Assortment Log

```

<appender name="AssortmentFileAppender" class="org.apache.log4j.RollingFileAppender">
  <param name="File" value="logs/.assortment.log"/><!-- = Runtime directory -->
  <param name="MaxFileSize" value="10240KB"/> <!-- = 10MB -->
  <param name="MaxBackupIndex" value="10"/>
  <param name="Threshold" value="TRACE"/>
  <layout class="org.apache.log4j.PatternLayout">
    <!-- The default pattern: DateTime Priority [Thread] [Category] Message\n -->
    <param name="ConversionPattern" value="%d{ISO8601} %-5p [%t] [%c{1}] %m%n"/>
    <!-- Pattern to output the caller's file name and line number -->
    <!-- param name="ConversionPattern" value="%d{ISO8601} %-5p [%t] (%F:%L) -
    %m%n"/-->
  </layout>
</appender>

<!-- Default logging level for assortment modifying and updating.
This category uses its own appender and does not route log events to other appenders.
Use "TRACE" to activate assortment loggers -->
<category name="assortment" additivity="false">
  <priority value="TRACE"/>
  <appender-ref ref="AssortmentFileAppender"/>
</category>

```

The enhanced logging will write detailed messages regarding assortment content calculation to *.assortment.log* file located in the server's logging folder. Please provide this log file to Informatica support.

2.4.5 Server and client in different time zones

At the moment it is only possible to use P360 server and client in different time zones with a max time offset of 23 hours.

Theoretically it is possible to have a server client constellation with a 26 hour offset. This is not supported. More than 23 hours offset ends up in exceptions like

```

java.lang.IllegalArgumentException: Timestamp format must be yyyy-mm-dd
hh:mm:ss[.fffffffffff]
    at java.sql.Timestamp.valueOf(Timestamp.java:237)
    ...


```

2.4.6 Cluster

Since 8.1.1.04 there is a file located at <sharedFolder>/shared/locks/datagrid/eventinfo.txt which shows the current state of the cluster of Product 360 servers.

This file is useful for trouble shooting fail over issues.

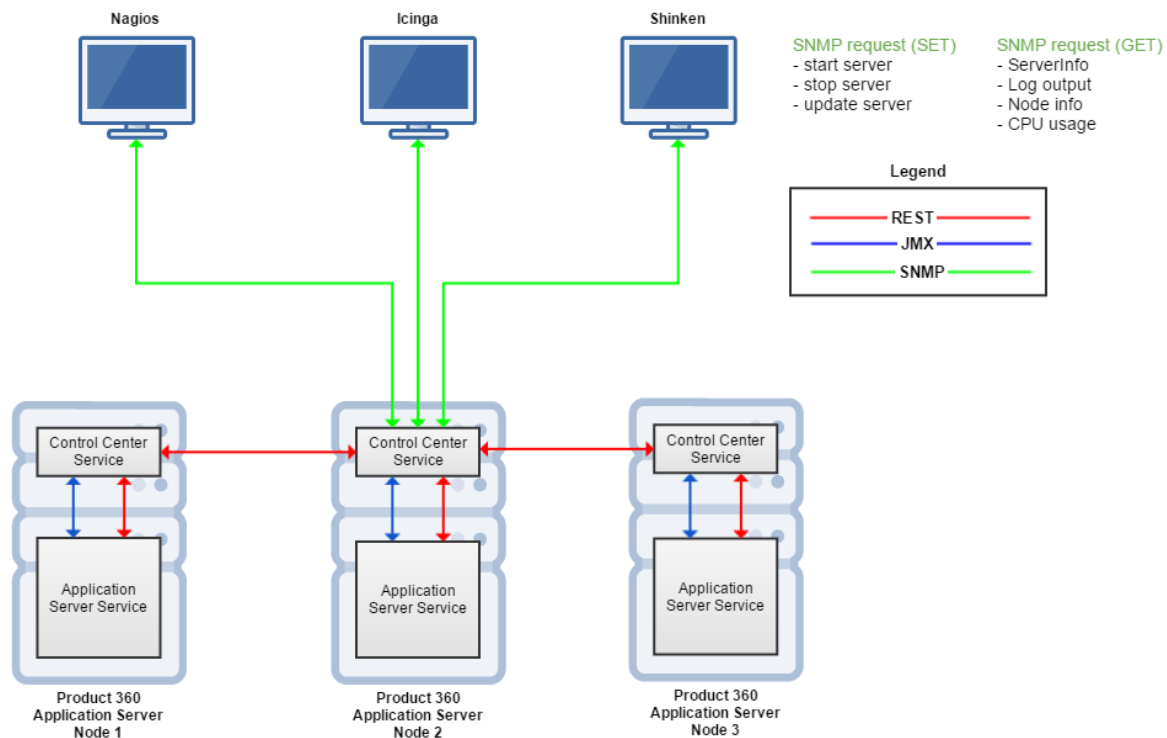
2.5 SNMP Monitoring

 SNMP Monitoring feature will be available with the version **8.0.03.00**

SNMP (Simple Network Monitoring Protocol) is an open standard for monitoring applications inside a bigger IT infrastructure. This page provides details as how to set up this monitoring feature for Informatica MDM - Product 360.

- [Overview](#)(see page 57)
- [Default Configurations](#)(see page 58)
 - [Agent](#)(see page 58)
 - [Network Management Station](#)(see page 58)
 - [Managed Objects](#)(see page 59)
 - [OID Configuration](#)(see page 59)
 - [Generation of OID configuration](#)(see page 59)
 - [Server Remote Command](#)(see page 59)
 - [Trap](#)(see page 60)
 - [Enable SNMP](#)(see page 60)
 - [Run application](#)(see page 61)
- [Modifying configurations](#)(see page 61)
 - [Change the agent's port](#)(see page 61)
 - [Network Management Station](#)(see page 61)
 - [Change the Trap receiver's port](#)(see page 61)
 - [Change the host address](#)(see page 61)
 - [Change user credentials](#)(see page 61)
 - [Add a user to the Network Management System configuration](#)(see page 62)
 - [Remove a user from the Network Management System configuration](#)(see page 62)
 - [Change OID Configuration](#)(see page 62)
 - [Server OID](#)(see page 62)
 - [Enterprise OID](#)(see page 62)
 - [Modifying Managed Object configuration](#)(see page 63)
 - [Enable/disable a managed object](#)(see page 63)
 - [Enable/disable notification for a managed object](#)(see page 63)
 - [Customizing notification for a managed object](#)(see page 63)
 - [Add notification to managed object](#)(see page 64)
 - [Advanced configuration](#)(see page 64)
 - [Multiple Monitoring](#)(see page 64)
 - [Add a new Managed objects](#)(see page 64)
 - [Add/remove managed Server](#)(see page 65)

2.5.1 Overview



The SNMP capable tools like Icinga, Nagios or others connect to a single Control Center service on any of the application server nodes. All control center services are connected with each other and therefore are able to provide the monitoring data for all the nodes in the Product 360 application server cluster. Communication between the Control Center services is based on REST, the communication between the Control Center service and its Application Server Service is done using REST and JMX.

2.5.2 Default Configurations

The system controller is delivered with standard configurations. Therefore these configurations can be modified in the configuration files.

2.5.2.1 Agent

- By default an SNMP agent is accessible through the port 161. This port number is also used by the system controller as default configuration.
- The system controller supports only communication using SNMP Version 3.
- The default protocol used by the system controller for communication through SNMP is UDP.
- Every notification is disabled by default except notification about the server status (STOPPED, STOPPING, and STARTING, RUNNING)

2.5.2.2 Network Management Station

- The control center provides by default configurations for a single Network Management Station (single monitoring). This station includes an SNMP manager in order to send SNMP request and receives traps. The

Network Management Station is located by default on the same host where the control center is running and traps are sent to the port 162.

- The default configuration of the Network Management Station offers 2 default users: one with a lower security level for sending GET request and receiving traps and the second user with a higher security level for sending both SET (start, stop and update server) and GET request and receiving traps. One of these users has the default username `controller`, the password `controller` and uses the authentication protocol SHA. The second user has the default username `Administrator`, the password `Administrator` and a private password `Administrator` and uses the authentication protocol SHA and the private protocol AES. To perform server remote control, users with this security level should be used.

2.5.2.3 Managed Objects

Every managed object is enabled by default to the managed server and a trap is sent once by events like Server STOPPED, STOPPING, STARTING, RUNNING. If you want traps to be sent by other events you should configure it (see section managed objects configuration).

2.5.2.4 OID Configuration

The generated OID configuration of each managed object is the result of the concatenation of the enterprise's OID, the server OID and the location of the object's value in the MIB. By default of the enterprise's OID is `.1.3.6.1.4.1.200`. The part `iso.org.dod.internet.private.enterprise(.1.3.6.1.4.1)` is the prefix. The number 200 represents the default enterprise identifier. Please change the 200 to your standard enterprise ID. If you don't have an enterprise identifier, please visit the link <http://pen.iana.org/pen/PenApplication.page> to request a unique identifier for your enterprise. Each managed server has an OID. If your environment has 2 servers, the first will have the OID 1.1 and the second 1.2. If you already have an SNMP Agent in your system that uses those OIDs, you should provide another to prevent MIB conflicts. (See section Modifying configuration)

2.5.2.5 Generation of OID configuration

At the program start a file is generated within the configuration directory with the name `"oid-data-config.properties"`. At the first successful connection of the system controller with each managed server a section is added to this file with the generated managed object OIDs. This operation is done for all managed servers. The generated OIDs are also reachable through SNMP by requesting the server description using the description OID. Each server has its description OID. It is found in the `oid-data-config.properties`.

2.5.2.6 Server Remote Command

All servers can be started, stopped or updated remotely using the remote command OID. To perform this operation you need a user that has the security level 2. The default remote command OID is `.1.3.6.1.4.1.200.3`. This OID is generated and is found in the server's description. The value expected for this object is a string that uses the following syntax: `command (start/stop/update) +", "+SERVER IDENTIFIER (e.g. pim-server1) +", "+MODE (kill, grace)`. The argument MODE is needed just to stop the operation. All given arguments should be separated by a semicolon. Errors in system controller during remote command operations are notified to Network Management Stations via traps. The trap OID of this error is a concatenation of the generated remote command OID and the sub string `.1`. If the generated

remote command OID is .1.3.6.1.4.1.200.3 the trap error Id will be .1.3.6.1.4.1.200.3.1. This information is useful for reading traps.

2.5.2.7 Trap

The trap is generated to announce an event. The generated traps by the Control Center uses a structure. It is a list of information about the event. The first element of this list represents the affected server and the second describes the event. The third element is the event level (Critical, Warning, and Ok) and the last the system time. The event level is used for management application to represent services with different colors (red, yellow and green) depending on the event level. For Monitoring Applications like Icinga or Nagios this level is represented by digits from 0 to 2 with following meanings:

- 0 = OK
- 1 = Warning
- 2 = Critical

Let us take an example. If at 7:55 the system detects that the server identified by PIM-SERVER1 is stopped. It will send the Trap with following information [PIM-SERVER1, STOPPED, 2, 7:55]. If 5 minutes after the stop the system detects that this is running (maybe the system administrator has start it after the event), a trap will be sent with following information [PIM-SERVER1, RUNNING, 0, 8:00]. Every trap has an OID. But all trap for the same event has the same OID. For each managed object the generated OID is used both for querying objects value via SNMP and for sending traps.

E.g. if the generated OID for the managed object `serverStatus` of the server identified by PIM-SERVER1 is .1.3.6.1.4.1.200.1.1.2.1 and the notification is enabled, all Trap notifications about this object will be sent using the same OID. The information about the structure of the trap sent and its OID are useful because you will need these for translation the trap and submit the result to the Network Management Station via passive check.

2.5.2.8 Enable SNMP

To enable SNMP, go to the configuration directory and find the file `NetworkConfig.xml`.

Edit this file as follow:

- Go to the xml element "`snmpdataconfigfile`"
- Remove the line which says : "`<!-- settings for snmp`"
- Also remove the line below the xml element "`snmp`", which says "`-->`"
- Save the file

Now find the `ClusterixConfig.xml` in the same directory and edit that file as follow:

- Again remove the line that says: "`<!-- settings for snmp`"
- Also remove the line at the bottom below the xml element "`NetworkManagementSystems`" that says "`-->`"
- Save the file and restart the application

2.5.2.9 Run application

After successful start of the Application, you should see the message Agent access `x.x.x.x:yyy` where `x.x.x.x` is the IP address of the host on which the agent is running and `yyy` the port number on which it is reachable. If this message appears, the system controller is ready for monitoring.

2.5.3 Modifying configurations

2.5.3.1 Change the agent's port

Use the VM argument `-Dclusterix.snmp.port` to change the agent's port. E.g. to set the port to 165 set the argument as follows `-Dclusterix.snmp.port=165`.

2.5.3.2 Network Management Station

Change the Trap receiver's port

1. Go to the configuration directory
2. Find the file `ClusterixConfig.xml`
3. Open it and go to the XML- root tag `NetworkManagementSystems` -> Inside this tag you will see a XML-tag with the name `NetworkManagementSystem`
4. Find his XML- element `<port>162</port>`
5. Change 162 to the port number you wish to assign.
6. Save the file and restart the application

Change the host address

1. Go to the configuration directory
2. Find the file `ClusterixConfig.xml`
3. Open it and go to the XML- root tag `NetworkManagementSystems`.
4. Inside this tag you will see a XML- tag with the name `NetworkManagementSystem`
5. Find his XML- element `<host>127.0.0.1</host>`
6. Change 127.0.0.1 to the IP address of your management station.
7. Save the file and restart the application.

Change user credentials

1. Go to the configuration directory
2. Find the file `ClusterixConfig.xml`
3. Open it and go to the your Network Management System definition
4. Go to the root element `Users`
5. Inside this tag you will see a XML- tag with the name `User`
6. Go to a user definition and edit his properties:
7. you want to change the username, find the XML-Element `<name>admin</name>`
8. Change admin to a preferred username.

If you want to change the password, find the XML-Element In the XML-Element `<password>Informatica</password>` Change **Informatica** to a preferred password. If you want to change the private password, and the XML element In the XML element `<privatePassword>Informatica</ privatePassword>` Change **Informatica** to your preferred private password.

At the moment, the private password is encrypted using the Advanced Encryption Standard with a 128 bit key.

Add a user to the Network Management System configuration

1. Go to the configuration directory
2. Find the file ClusterixConfig.xml
3. Open it and go to the your Network Management System definition Go to the XML- root Users
4. Inside this tag you should find a XML- tag with the name User
5. Duplicate a user definition and edit the properties as explained above. Leave the XML element `privatePassword`, and `privateProtocol` blank (`<privatePassword> </ privatePassword>`, `<privateProtocol></ privateProtocol>`) if the user you want to add belongs to the security level 1.
6. Save the file and restart the application.

Remove a user from the Network Management System configuration

1. Go to the configuration directory
2. Find the file ClusterixConfig.xml
3. Open it and go to the your Network Management System definition
4. Go to the XML - root Users
5. Inside this tags you will see a XML- tag with the name User
6. Remove the tag of the user you want to delete to your system.
7. Save the file and restart the application. Please note that the SNMP Agent needs at least one user with security level 2 to start.

2.5.3.3 Change OID Configuration

Server OID

1. 1. Go to the configuration directory
2. 2. Find the file NetworkConfig.xml
3. 3. This file contains the Network configuration of all managed server. Each managed server is represented by an XML-tag node.
4. 4. Find the node of the server which the configuration should be edited.
5. 5. Find the XML-Element `<snmp oid=1.1>`
6. 6. Replace 1.1 with the Id you want. Please note that the Id consists of separated numbers.
7. 7. Save the file and restart the application.

Enterprise OID

The enterprise OID entry is found in the Object configuration file of each server.

1. Go to the configuration directory
2. Find the file NetworkConfig.xml .This file contains the Network configuration of all managed server. Each managed server is represented by the XML-tag node.
3. Find the node-definition of the server which the configuration should be edited.
4. Find the XML-Element `<snmpdataconfigfile name="file name">`

5. Check the assigned file name and remember it
6. Look this file in the configuration directory and open it
7. Go to the XML-tag general
8. Edit the XML-element enterprise(change the last number of the property id (200) to your enterprise ID)
9. Save the file and restart the application.

Modifying Managed Object configuration

1. Go to the configuration directory and find the file NetworkConfig.xml
2. This file contains the Network configuration of all managed server. Each managed server is represented by the xml-tag node.
3. Find the node-definition of the server for which configurations are to be edited.
4. Find the xml-Element <snmpdataconfigfile name=file name>
5. Check the assigned file name and remember it
6. Look for the file in the configurations directory and open it
7. Go to the XML-tag rest for rest managed objects and jmx for jmx managed objects:

Enable/disable a managed object

1. Check the property append
2. Set the value of this property to true to enable the object or false to disable the object.

Attention: Enabling/disabling objects is a critical operation because it affects all your monitoring system configurations (your management application and the system controller) if you already had a running monitoring solution. After this operation all or some managed objects become a new OID depending on their position in the oid-config-file. Please compare the new generated OID of each managed object with the OID you previously configured in your monitoring application to be sure that it is the same. If it is not the same use the new generated OID of the object. Otherwise you can have the following problems/errors in your monitoring application:

- Some OID are not found (If you disabled managed objects)
- Some OID are found but their value belong to another objects.

Enable/disable notification for a managed object

This works if the property append is set to true

1. Check the property notification_enabled
2. Set the value of this property to true to enable notification or false to disable the notification.

Customizing notification for a managed object

This works if the properties append and notification are set to true.

- If the property type equals to String, assign a value to the property pattern. This represents the value you expected for this managed object. If another value is found over the time, a Trap-notification will be send to all register Network Management Station (see section Network Management Station above) with the actual value of the object.
- If the value of the property type different to String (e.g. int or long).Please assign a range to the property range This represents the possible values you expect for the managed object over the time.

Note: if a value is found over the time, that do no belong to the range, a notification will be send register Network Management Stations with the actual value of the object. To give the range follow the syntax [+min value+.. +max value +]. The bracket [] means that the minimal and the maximal value belong to the range. Attention: if you have a syntax error during range configuration or the value you assigned to min is greater than the value of max the system controller will start but the configured notification will be ignored. Please check the log file for error description.

Add notification to managed object

You can add notifications to every managed object you want. But be sure that append is set to true.

- To add a notification, add the following property to the managed object: `notification_enabled=true`.
- If this object has the type String adds the property `pattern=`The expected value of this object otherwise adds the property `rang= [min...max]` (see section customizing notification for a managed object above).

2.5.3.4 Advanced configuration

Multiple Monitoring

The standard configuration allows the system controller to send trap just to one Network Management Station. To enabled multiple monitoring do the following:

1. Go to the configuration directory
2. Find the file `ClusterixConfig.xml`
3. Open it and go to the XML- root `NetworkManagementSystems`. In this tag you will see a XML- tag with the name `NetworkManagementSystem`
4. Duplicate this tag and edit like explained above (see section Network Management Station).
5. Save the file and restart the application. Every Trap will be sent to all Network Management Stations listed in this configuration file. If a trap is send to a Network Management Station and a acknowledge is not received, the trap will be send twice. If the second trap is lost, an error will be log.

Add a new Managed objects

This works just for the JMX communication.

1. Go to the configuration directory and find the file `NetworkConfig.xml`. This file contains the Network configuration of all managed server. Each managed server is represented by the xml-tag node.
2. Find the node-definition of the server which the configuration should be edited.
3. Find the xml element `<snmpdataconfigfile name=file name>`
4. Look the assigned file name and remember it
5. Look this file in the configuration directory and open it
6. Go to the XML-tag JMX
7. Duplicate a JMX-Object definition
8. Edit as follow
 - Object name: set the MBean Object Name. You can use the JConsole Widows (see <http://docs.oracle.com/javase/tutorial/jmx/remote/jconsole.html>) to have the right name of the MBean Object.
 - Element name: set the attribute name of the MBean Object to query.
 - Set the type of this attribute as showed in the JConsole Windows. Note: If the type is String you cannot have the graphic visualization of this object in your Network Management Station. It is done just for performance data.

- Set access to read_only or read_write depending on your choice. Note that if you set the access type to read_write the value of this object can be change via SNMP-Set request. And you should need specials rights for this.
- Set append to true to enable the managed object.
- Add notification as explained in the section "Add notification" above if you need any
- Save the file and restart the application.

Note: If the given MBean object name of the managed object that you try to add or the MBean attribute you query does not exists or is not correctly spelled, the managed object will be ignored and an error will be log.

Add/remove managed Server

1. Go to the configuration directory and find the file NetworkConfig.xml. This file contains the Network configuration of all managed server. Each managed server is represented by a xml-tag node.
2. Duplicate a node-definition and edit it as follow:
 - Go to his xml element JMX and give the port number of this server
 - Go to his XML-element OID and give an SNMP OID to this server (e.g 1.3)
 - Go to the XML-element snmpdataconfigfile and replace the file name with snmp-data3.xml
 - Duplicate a snmp-data-denition file and rename it (e.g snmp-data3.xml)
 - Configure managed objects as explained above
 - Save the file and restart the application.

A new managed server is added to your system controller and you can query the value of the managed object and receiving traps using the generated OIDs.

2.6 DB Connection pool

PIM Core database workload consists of a large amount of parallel requests. To handle such workload PIM Core is using DB connection pool. Connection pool holds established jdbc connections and returns them to the pool when higher level data access layer such as Hibernate or UDA close them. Such pool helps to save time required to establish and prepare jdbc connection. However connection pool increase memory consumption. PIM Core maintains one connection pool per data source (MAIN,MASTER,SUPPLIER) and per data access module (UDA,Hibernate), total 6 connection pools. Quartz and JBPM maintain their own connection pools. Due to performance reasons PIM Core is using tomcatjdbc connection pool which is optimized for concurrent access. Popular implementations DBCP and c3p0 are not suitable for PIM Core workloads.

It is important to configure connection pool appropriately. See <https://tomcat.apache.org/tomcat-8.0-doc/jdbc-pool.html> for more detail on tomcat jdbcpool configuration options.

To improve jdbc performance the connection pool is capable of configuring prepared statement caches. This means prepared statements are cached per connection, and please note that this will significantly increase memory consumption! Also the connections need to be configured in a way that they do not remove idling connections too fast otherwise you will render the cache useless.

✓ JMX

Use hibernate and connection pool JMX managed beans to analyse workload and adjust values on production systems, see also [Tuning Advisory](#)(see page 65)

2.6.1 High performance connection pool configuration example

The default configuration should be sufficient for regular workloads. However in context of having many concurrent users and/or system jobs (import, export, data quality, etc.) the connection pool might become a bottle neck. The connection pool will throw an error in case the connections are exhausted, but this will not necessarily happen because there is a wait time until the error will be thrown. To adjust connection pool settings, the following data access template files must be adjusted.



Oracle

When running PIM on Oracle do not enable caching of prepared statements unless you make sure that initialization parameter **OPEN_CURSOR** is set high enough. The limit is counted globally for all connections that belong to the same pool.

2.6.1.1 Hibernate

The template files when running PIM on Oracle for the Hibernate connection pools are located in the **server\configuration\HPM\database** folder:

- main.properties.template.ORA11g
- master.properties.template.ORA11g
- supplier.properties.template.ORA11g

In case you are running PIM on Microsoft SQL Server, the corresponding template files have the MSSQL2008 suffix.

Make sure to modify the template files since PIM will always generate the real property files during server startup.

✓ Configuration for 250 maximum active connections

increase validation interval (value in milliseconds) to avoid permanent connection check which is unnecessary in today's networks
hibernate.tomcatjdbc.pool.validationInterval = 120000

increase eviction interval (value in milliseconds) when the pool should check which idling connections could be removed from the pool
hibernate.tomcatjdbc.pool.timeBetweenEvictionRunsMillis = 30000

increase idle time (value in milliseconds) before the connection may be removed from the pool
hibernate.tomcatjdbc.pool.minEvictableIdleTimeMillis = 1800000

increase the amount of initial connections since creating new connections is a very time intensive task
hibernate.tomcatjdbc.pool.initialSize = 50

increase the amount if minimum idling connections which will be kept open at any time to fully utilize statement cache and avoid hard parsing repeating queries
hibernate.tomcatjdbc.pool.minIdle = 125

increase the amount of maximum idling connections to avoid closing idling connections during peak times
hibernate.tomcatjdbc.pool.maxIdle = 250

increase the amount of maximum idling connections to avoid the situation that PIM processes are starving from available connections
hibernate.tomcatjdbc.pool.maxActive = 250

lower the maximum wait time (value in milliseconds) until the connection pool will throw an exception that the pool is exhausted
hibernate.tomcatjdbc.pool.maxWait = 10000

enable jdbc interceptor for statement cache, increase max value if required by application load pattern
hibernate.tomcatjdbc.pool.jdbcInterceptors =
ConnectionState;StatementFinalizer;StatementCache(prepared=true,callable=false,max=250)

```
# Oracle only: this will force the jdbc driver to enable its builtin
statement cache to avoid hard parsing repeating queries
hibernate.tomcatjdbc.pool.connectionProperties =
oracle.jdbc.implicitStatementCacheSize=250

# Oracle only: prepared statements (INSERT or UPDATE) can be batched if
batch_size is > 1
hibernate.jdbc.batch_size = 50
```

All other settings should be left untouched. Increasing the amount of connections and/or statement cache size will result in higher memory consumption, so make sure that application server and database memory target is sized properly.

2.6.1.2 UDA

The template file when running PIM on Oracle for the UDA connection pools is located in the **server\configuration\HPM\database\uda** folder:

- `uda-jdbcpool.xml.template.ORA11g`

In case you are running PIM on Microsoft SQL Server, the corresponding template file has the MSSQL2008 suffix.

Make sure to modify the template file since PIM will always generate the real property files during server startup.

This example will adjust the same values like the Hibernate connection pool (please note that this is a xml file):

✓ Configuration for 250 maximum active connections

```

<!-- increase validation interval (value in milliseconds) to avoid
permanent connection check which is unnecessary in today's networks -->
<property name="validationInterval">120000</property>

<!-- increase eviction interval (value in milliseconds) when the pool
should check which idling connections could be removed from the pool -->
<property name="timeBetweenEvictionRunsMillis">30000</property>

<!-- increase idle time (value in milliseconds) before the connection may
be removed from the pool -->
<property name="minEvictableIdleTimeMillis">1800000</property>

<!-- increase the amount of initial connections since creating new
connections is a very time intensive task -->
<property name="initialSize">50</property>

<!-- increase the amount if minimum idling connections which will be kept
open at any time to fully utilize statement cache and avoid hard parsing
repeating queries -->
<property name="minIdle">125</property>

<!-- increase the amount of maximum idling connections to avoid closing
idling connections during peak times -->
<property name="maxIdle">250</property>

<!-- increase the amount of maximum idling connections to avoid the
situation that PIM processes are starving from available connections -->
<property name="maxActive">250</property>

<!-- lower the maximum wait time (value in milliseconds) until the
connection pool will throw an exception that the pool is exhausted -->
<property name="maxWait">10000</property>

<!-- enable jdbc interceptor for statement cache, increase max value if
required by application load pattern -->
<property
name="jdbcInterceptors">ConnectionState;StatementFinalizer;StatementCache
(prepared=true,callable=false,max=250)</property>

<!-- Oracle only: this will force the jdbc driver to enable its builtin
statement cache to avoid hard parsing repeating queries -->
<property
name="connectionProperties">oracle.jdbc.implicitStatementCacheSize=250</
property>

```

Again, all other settings should be left untouched. Increasing the amount of connections and/or statement cache size will result in higher memory consumption, so make sure that application server and database memory target is sized properly.

2.7 Server Job Maintenance

Server jobs store their current state as well as all protocol entries in the database. A special job is executed once a day and is able to delete other, finished, jobs.


- [Configure the life time of a job](#)(see page 70)
- [Syntax](#)(see page 70)
 - [Example](#)(see page 70)
- [Standard Job List](#)(see page 70)

2.7.1 Configure the life time of a job

A job's life time is either defined as default life time by the job contribution itself (`plugin.xml`), or with the `plugin_customization.ini` file.

The values defined in the `plugin_customization.ini` will always have precedence over the default of the contribution.

To define a life time you need the unique job identifier (as listed in the table below), e.g. for the merge it's "Merge".

 A lifetime of 0 (zero) means no cleanup of the job, thus these jobs will not be removed from the job-history.

2.7.2 Syntax

```
com.heiler.ppm.jobhistory.cleanup.server/{JobTypeIdentifier}.lifeTime = {amount in days}
```

2.7.2.1 Example

```
com.heiler.ppm.jobhistory.cleanup.server/Merge.lifeTime = 60
```

Means that already executed merge jobs are deleted after 60 days of their execution.

2.7.3 Standard Job List

List of standard jobs with their identifiers in order to configure their lifetime.

Job identifier	Job name	Job group	Default life time	Multiple Job entries
AssignDocumentJob	Media manager document assignment	System processes	3	TRUE
CheckTaskEscalatedJob	Check for escalated tasks	System processes	0	TRUE
CleanUpJobHistory	Remove obsolete processes	System processes	0	TRUE
CleanUpReports	Remove obsolete query results	System processes	0	FALSE
CleanUpSoftdelete	Remove deleted objects	System processes	0	FALSE
com.heiler.ppm.revision.jobType	Versioning	Data maintenance	10	TRUE
com.heiler.ppm.workflow.server.CleanUpFinished	Remove obsolete workflows	System processes	0	FALSE
CopyValueIntoStructureGroupsJob	Update structure group features	System processes	0	TRUE
create.thumbnails	Thumbnail generation	System processes	3	TRUE
create.thumbnails.rec	Recursive thumbnail generation	System processes	0	TRUE
DataQualitySeries	Data quality check (scheduled)	Data quality	0	TRUE
DictionarySynchronization	Dictionary synchronization (manual)	Data quality	0	TRUE
ExportDryRun	Quality check	Data maintenance	0	FALSE
ExportSeries	Export operations (repeated)	Publications	0	TRUE
FindReplace	Find and Replace	Data maintenance	2	TRUE

Job identifier	Job name	Job group	Default life time	Multiple Job entries
FindReplaceSetValue	Set	Data maintenance	2	TRUE
FulltextSearchJobSeries	Index creation (repeated)	Search index	30	TRUE
FulltextSearchSingleJob	Index creation (one-off)	Search index	30	TRUE
Import	Import operations	Data transfer	0	TRUE
Inbox	Hotfolder data processing	Data transfer	0	TRUE
InboxImport	Import operations by hotfolder	Data transfer	0	TRUE
ManualCleanUpReports	Remove obsolete query results (manual)	System processes	0	FALSE
MapArticlesToProductGroupJob	Product item feature transfer	System processes	0	TRUE
MapArticlesToStructureGroupJob	Update item-structure group assignments	System processes	0	TRUE
MapEGDsToStructureGroupJob	Update EGD-structure group assignments	System processes	0	TRUE
Merge	Merge operations	Data transfer	0	TRUE
MigrateStructures	Migrate Structures	Data maintenance	0	TRUE
NewMediaAssetDerivativeSchema	New derivative schema	System processes	0	TRUE
ReferenceDataDeployment	Deployment of reference data (scheduled)	Data quality	0	TRUE
SingleDataQuality	Data quality check (manual)	Data quality	0	FALSE
SingleExport	Export operations (one-off)	Publications	0	FALSE

Job identifier	Job name	Job group	Default life time	Multiple Job entries
SupplierExchangeExport	Export operations (Supplier Portal)	Publications	5	TRUE
TestExport	Test export operations	Publications	0	FALSE
update.version	Updating image versions	System processes	0	TRUE
UpdateTasksJob	Update tasks	System processes	0	TRUE

2.8 Soft Delete Cleanup Job

2.8.1 Why is there a concept of soft delete and a cleanup job that triggers the logging in the DB?

The soft delete mechanism allows to share information on deleted records in our DB with external systems using e.g. the deleted data providers in the Export. However, if you would keep all soft deleted objects in the DB forever it would lead to a DB space explosion and severe performance stretch and therefore a job can be configured to clean up soft deleted records from the DB after X days. A good thing to have for an enterprise application if you have a use case for this. If not, you may also want to switch the soft delete of completely and save some space on your DB.

Change the configuration in "plugin_customization.ini" and restart service

```
# Defines the default deletion mode in case entities support both, soft
and hard delete

# Supported values are SOFT and HARD

com.heiler.ppm.std.server/default.deletion-mode = HARD
```

Note: If the recovery logging level of your DB is set to FULL any change is stored in the transaction log for a roll back in case of errors. Otherwise you could only roll back using DB backups.

2.8.2 General information

The soft delete cleanup job is a maintenance job which performs a hard delete (physical delete) of entities which have been already soft deleted. The hard delete cleanup is only triggered for expired entities and only if the lifetime is not unlimited. With version higher 8.0.01, the lifetime is set to one year per default for all entities. The expiration date is calculated by means of the current date as basis. So if today is February the 5th and you have a lifetime of 1 month, then the calculated expire date will be January the 5th, which means

that all entities deleted prior to this date will be removed permanently. The default life time property defines life time for all entity types. This can be overwritten per root and sub entity type.

2.8.3 Configuration

You can define the cleanup life time in the *plugin_customization.ini*. Search for `cleanup.lifeTime.default`.

Example default life times for all entities:

```
cleanup.lifeTime.default=1d
cleanup.lifeTime.default=1y
```

If you want to define a different life time for different root and sub entities then you need to uncomment the property `cleanup.lifeTime.custom.file` and refer to a custom file defining the soft delete properties. The custom file has to be located in the folder `~/server/configuration/HPM`. The `cleanup.lifeTime.custom.file` property referring to the custom file could look like this, for example:

```
cleanup.lifeTime.custom.file = softDeleteLifeTime.properties
```

In your customized soft delete life time properties file you can define the different life times for the entities. The most concrete definition will be used. So if you define a life time for a root entity, then you can overwrite it for the sub entity. But this is only valid if the life time for the sub entity is shorter than the one of the parent.

A valid configuration:

```
cleanup.lifeTime.default = unlimited
ArticleType.lifeTime = 1y
ArticleLangType.lifeTime = 1m
StructureGroupLangType = 6d
```

2.8.4 Schedule

To schedule the job repetitively, uncomment the property `cleanup.job.repeatPattern` and define the pattern. From the version 8.0.01.00 this job is scheduled to run every day at 24:00. To disable the scheduling, you have to uncomment the property and leave it empty, then the job won't be executed automatically but may be started manually in the process overview (Select "System processes > Remove deleted objects", select the process in the "Current and past processes" view and call "Execute immediately" from the context menu).

Example for running the job daily at 24:00:

```
cleanup.job.repeatPattern = 0 00 00 * * ? *
```

To disable scheduling:

```
cleanup.job.repeatPattern =
```

For more information refer to the Quartz Enterprise Job Scheduler documentation for Cron Expressions

2.8.5 Be cautious with mass deletes in such a setup

The recovery level FULL would increase your transaction log if your DBA is not securing the log away regularly. Potentially endless so to speak, or just until the disk space runs full. A one-time mass deletion of records in Product 360 under recovery level FULL may cause such problem once the cleanup job kicks in.

Recovery level SIMPLE for example would delete the log after every commit to free up log space again.

2.8.5.1 How do other customers solve such a constellation?

So far, we have not seen many customers deleting that many records in such a short period of time. However, we have seen 4 different variations on best practices in this scenario:


1. They give the DB as much space as it would need. Even temporary...
2. They have the cleanup job executed frequently in general (typically delete everything older than 4 weeks) to keep the overall amount of soft deleted records as low as possible over time
3. In case of a one-time mass deletion process customers create a DB backup and then set the logging level to SIMPLE for the time needed by the transaction
4. They back up the transaction log prior this action to start with an empty log and keep disk space consumption as limited as possible

These are scenarios where it is already identified that this mass deletion could otherwise create a problem. However, what happens if you didn't know about this and now are stuck with a transaction log running full all the time the cleanup job gets executed?

2.8.5.2 Possible Workarounds

Set the X days a record should be kept in the DB relatively high to reduce the amount of overall records in the queue to be cleaned up as small as possible per job execution. This needs some further monitoring over time but is the best option in a safe setup to reduce the amount of soft deleted records controlled and as granular as possible.

The key question is whether you have a lot of soft deleted records in the queue (e.g. because a clean-up job has only recently been configured for the first time in your setup) or if that space problem is caused only having deleted records in the queue that share the same day on their deletion time stamp. If the latter is the case you could still try to work with the bulk size property to get things solved by defining the number of records to be taken by the job per transaction:

 Add this line to the plugin_customization.ini of the server to limit the number of items per transaction to 1000 (or whatever suits your idea, default is 100,000). Better use control center to adjust the file and deploy it on the servers if you have multiple servers.

```
com.heiler.ppm.std.server/entity-manager.delete.bulk-size = 1000
```

2.9 Data Quality Operation

2.9.1 Logging

Additional logging information for data quality execution can be increased on several levels: rule engine execution, data provisioning and performance analysis

2.9.1.1 Rule engine execution

In order to see logging information of the execution of a data quality rule by the rule engine set the category 'IDQ' to DEBUG in server log4j.xml

```
<category name="IDQ">
  <priority value="DEBUG"/>
</category>
```

As opposed to normal logging information a server restart is required for changes to take effect. The logs can be found in separate files in the server logs folder. Each of them has a pattern like <Execution Thread name>-<Executed rule name>-<Execution start time>.log

2.9.1.2 Data provisioning

In order to see logging information of the data provisioning (data input into and data output from rule) set the category 'DQ_EXECUTION' to DEBUG in server log4j.xml.

```
<category name="DQ_EXECUTION">
  <priority value="DEBUG"/>
</category>
```

The output is available in the usual error.log respectively console.log files. Example output:

```

13:01:25,166 DEBUG [exe-IDQ-SDK-Check_PresetValues-2] [DQ_EXECUTION] DQ IN line: 0
(rule: preset violation)
  inObjectID: '33762@1' (field: Object_ID)
  inAttributeName: 'Color' (field: ArticleAttribute:name)
  inLanguage: 'English' (field: ArticleAttributeLang:language)
  inStructure: '10058' (field: ArticleAttributeStructureGroupAttributeMap:structure)
  inStructureGroup: '232@10058' (field:
ArticleAttributeStructureGroupAttributeMap:structureGroup)
  inStructureFeature: '74@10058' (field:
ArticleAttributeStructureGroupAttributeMap.StructureAttribute)
  inAttributeValue: 'green2' (field: ArticleAttributeValue.Value)
  inValueIdentifier: 'DEFAULT' (field: ArticleAttributeValue:identifier)
  inStructurePresetValue: '' (field:
ArticleAttributeStructureGroupAttributeMap.DomainValue)
  inDataType: 'Character string' (field: ArticleAttribute.Datatype)

13:01:25,317 DEBUG [exe-IDQ-SDK-Check_PresetValues-1] [DQ_EXECUTION] DQ OUT line: 0
(rule: preset violation)
  status_code: '0' (field: QualityStatusEntry.Status)
  status_message: 'Primary, Pflichtgruppe, Feature 'Color': Value 'green2' (English,
DEFAULT) of attribute 'Color' does not conform to the preset values' (field:
QualityStatusEntry.Message)
  object_id: '33762@1' (field: Object_ID)

```

2.9.1.3 Performance analyzing

Enable the following category to DEBUG for additional information about DQ initialization parameters and DQ Executor Pool usage.

```
<category name="com.heiler.ppm.dataquality.server">
  <priority value="DEBUG"/>
</category>
```

2.9.2 Performance tuning

- Increase number of maximum jobs executed in parallel by adjusting the setting **com.heiler.ppm.job.server/maxRunningServerJobs** accordingly.
- Each pool properties are available via JMX and also the size of each DQ Executor Pool can be adjusted at runtime via JMX (MBeans → com.heiler.ppm → dataQuality → dqExecutorPools)

[illegible]

- Use synchronous execution of DQ via REST API instead of asynchronous jobs and waiting for job finished result (e.g. when executing rules inside a bpm workflow).
This decreases the general load of the job framework and avoids the general overhead of job scheduling; see also: Rest Management API -> Rest Data Quality API -> Rule Execution.
- The precision of the mapping input ports should be sized as small as reasonably possible, since this has a direct impact on rule execution performance and DTM memory allocation
- Rule execution performance could benefit from setting a lower value for **com.heiler.ppm.dataquality.server/dataquality.executor.pageSize** as the data sets that are needed to be retrieved will be smaller, hence faster
- Also consider having a smaller value for **com.heiler.ppm.dataquality.server/dataquality.executor.pageSize** in general since the executor can then be used earlier in a different execution scenario while the system is busy writing the result of the current elements of the DQ Execution Result to the database (the executor pool serves FIFO)
- Rule execution performance could benefit from setting a higher value for **com.heiler.ppm.dataquality.server/dataquality.execution.dtmMaxMemPerRequest** to avoid that DTM has to swap out cache and index files to disk during rule execution, however this is subject to the available physical memory of the host. An alternative could be moving the temporary DQ folder to a RAMDisk by setting **com.heiler.ppm.dataquality.server/dataquality.execution.tmpFileLocation**
- Depending on the hardware the setting **com.heiler.ppm.dataquality.server/dataquality.execution.maxThreads** can be increased to utilize the available CPU cores

The following properties in the preferences.ini show default values and may be adjusted to improve performance depending on corresponding use-cases:

```

# Enables an alternative implementation for the DQ Rule Check_IsEmpty that doesn't
use the IDQ SDK.
# Set to true in case you want to avoid DTM overhead for simple isEmpty checks.
# Default: false
com.heiler.ppm.dataquality.server/dataquality.alternativeCheckIsEmpty=false

# Specifies the page size when loading the list model of data to be checked for IDQ
rule execution.
# Default: 10000 objects
com.heiler.ppm.dataquality.server/dataquality.executor.pageSize=10000

# Maximum amount of concurrent executors in the executor pool.
# Each executor can run independent from each other to avoid contention on specific
IDQ rules.
# Default: 2
com.heiler.ppm.dataquality.server/dataquality.executionPool.maxTotalPerRule=2

# Example: maximum amount of concurrent executors in the executor pool per specific
IDQ rule.
# The '/' in the rule identifier has to be replaced with '_'
com.heiler.ppm.dataquality.server/
dataquality
.executionPool.maxTotalPerRule.Informatica_PIM_Content_PreDefined_Rules_Check_DataTyp
es=5

# Specifies in seconds how long the MappingExecutor instances will be kept in the
pool before aging out.
# Exceptions:
# - Executors will be destroyed immediately in case a new version of a rule will be
uploaded
# - Executors will be destroyed immediately when related dictionaries have been
changed
# - Executors can only age out without a dictionary change in case they have been
used at least once
# Default: 3600 seconds (60 minutes)
com.heiler.ppm.dataquality.server/
dataquality.executionPool.executorKeepAliveInterval=3600

# Example: specifies in seconds how long the MappingExecutor instances will be kept
in the pool before aging out per specific IDQ rule.
# The '/' in the rule identifier has to be replaced with '_'
com.heiler.ppm.dataquality.server/
dataquality
.executionPool.executorKeepAliveInterval.Informatica_PIM_Content_PreDefined_Rules_Che
ck_DataTypes=14400

# Specifies the amount in bytes the DTM should be allocating from the host OS memory
(note that this is not JVM heap memory).
# If configured properly this will avoid swapping cache files to disk during
transformations.
# Default: empty (DTM will automatically determine a value)

```

```

dataquality.execution.dtmMaxMemPerRequest=

# Specifies the amount in bytes the DTM should be allocating from the host OS memory
# per specific IDQ rule.
# The '/' in the rule identifier has to be replaced with '_'
#dataquality.execution.dtmMaxMemPerRequest.Informatica_PIM_Content_PreDefined_Rules_C
#heck_DataTypes=58720256

# Specifies the amount of threads in the IDQ SDK executor thread pool resp. how many
# executors can run at the same time.
# This roughly translates to dedicated CPU cores for IDQ execution.
# Default: 5
com.heiler.ppm.dataquality.server/dataquality.execution.maxThreads=5

# Specifies the location for the temporary files for each executor that are created
# during IDQ execution.
# Default: empty (using 'workspace/.metadata/com.heiler.ppm.dataquality.server' as
# default)
dataquality.execution.tmpFileLocation=

```

An additional note to the property **dataquality.executionPool.maxTotalPerRule**: all rule executors will be already created during server startup in background threads.

2.9.2.1 Cleanup of Status Entries

When executing data quality, the execution result is written to the database table `StatusEntry`, referring to the used executed data quality rule configuration.

If a data quality rule configuration is deleted (e.g. via Desktop Client DQ Perspective), the status entries for that deleted configuration are not deleted immediately.

A cleanup job will handle this, which runs every day at midnight per default.

The identifier of that job is 'CleanUpDataQualityStatusEntries' and the (english) name 'Status entries cleanup (scheduled)'.

plugin_customization.ini

```

com.heiler.ppm.dataquality.server/dataquality.status.cleanup.job.pattern = 0 0 0 ? *
1,2,3,4,5,6,7

```

Additionally, it is recommended to activate soft deletion of status entries. To do this, the already existing soft delete job has to be used.

Please refer to the 'Soft Delete Cleanup Job' chapter in the 'Server Operation' documentation and apply it to the sub entity type 'StatusEntryType'. Also ensure, to configure the lifetime for the root entity type 'StatusType' to another value than "unlimited" (e.g. '1y'), and be aware, that thereby also other sub entity types of 'StatusType' (e.g. 'PublicationStatusEntryType', which is used by the GDSN Accelerator) will also be cleaned up.

This is because the most concrete definition for an entity will be used. So if you define a life time for a root entity, then you can overwrite it for the sub entity. But this is only valid if the life time for the sub entity is shorter than the one of the parent.

2.9.3 Issue reporting guide

This guide explains required information and steps to retrieve them in order to ensure a certain overall quality level for issues concerning Data Quality relevant topics.

Please provide as much information as described as follows in order to have a quick success in solving the issues.

2.9.3.1 Configuration setup

All DQ relevant configuration and rule files contained in the server configuration subfolder 'dataquality'. This can simply be done by zipping the whole folder and attaching it to the issue.

Screenshots of specific rule configurations are also helpful, especially when they are custom ones. The same is true for the configured execution triggers.

It may also important to provide information about the setup workflows if executions are called from there.

2.9.3.2 Information about the execution context

Data Quality executions have following main informations

- The rule configurations containing information about the data provisioning
- The actual data quality rules that are used in the rule configuration
- The execution may contain more than one rule configuration, mostly if the execution is done for specific channels containing several configurations
 - If the rule configurations are executed as jobs as f.g. by manual executions, see the job log details in the processes overview for the used rule configurations.
 - For triggered executions this information may also be retrieved by looking at the trigger configuration
 - Execution may be also triggered from calls inside workflows
- Can the problem be pinned down to a specific rule configuration or rule? If so, Please specify them. Is the rule a custom rule and if so, what is it supposed to do?

2.9.3.3 Logging data

- General logging data as always (error.log, console.log)
- DQ specific logging data as described in the previous 'Logging' section. This may require to change the logging level to DEBUG as described there and then execute the rules again in order to see the additional logging information.

2.9.3.4 Performance

- Provide information about the settings as described in the Performance tuning section
- Is the execution manually or triggered e.g. after import or every night or for each item change?
- How many items are affected? Execution for single item or many items of a catalog or other data set f.g.?
- Are there many executions in parallel?

2.9.3.5 Crashes

In case of crashes get the heap and thread dump of the P360 Server.

If the crash happens on the DQ engine itself, also provide heap dumps of the underlying engine shared libraries.

For Windows this can be done like this:

- ensure
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\Windows Error Reporting\LocalDumps
registry key is present
- try to reproduce the problem. Dumps are available in
%LOCALAPPDATA%\CrashDumps
or
C:\Windows\System32\config\systemprofile\AppData\Local\CrashDumps
in case service is run with system account.

2.9.4 Configure rules that are affected by Dictionary Synchronization in order to increase performance of IDQ rule execution

2.9.4.1 Situation/ motivation

- For each IDQ rule, PIM creates executor objects and stores them in a pool, because creation is expensive. Executors are used to execute IDQ rules in parallel. These executors are created lazily on demand. Also IDQ SDK internally synchronizes executor creation, which means, creation is queued.
- Changes to dictionaries aka reference tables are only picked up by the rule if the executor is disposed and re-created
- Currently, on the PIM side, it is not known which reference tables are used by which IDQ rule. Hence, whenever, a dictionary is synchronized, all executors are disposed, and created lazily on demand.
- Because of that the performance drastically degrades whenever a dictionary is synchronized AND many new executors are requested for IDQ rule execution.

2.9.4.2 Solution

- It is now possible, to configure, which IDQ rules are affected by which dictionary (aka Reference Table). Therefor the **DictionaryRuleConfiguration.xml** is used, which is stored in the PIM server's folder / configuration/HPM/dataquality, and looks like this:

DictionaryRuleConfiguration.xml

```
<?xml version="1.0" encoding="UTF-8"?>
<dictionaryToRuleConfiguration xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="DictionaryRuleConfiguration.xsd">
  <isActive>true</isActive>
  <dictionary>
    <dictionaryIdentifier>Informatica_DQ_Content\Dictionaries\General\profanity_infa.
dic</dictionaryIdentifier>
    <affectedRules>
      <ruleIdentifier>Informatica_PIM_Content/PreDefined_Rules/Check_Profanity</
ruleIdentifier>
      <ruleIdentifier>Informatica_PIM_Content/Custom_Rules/MyCustomRule1</
ruleIdentifier>
    </affectedRules>
  </dictionary>
</dictionaryToRuleConfiguration>
```

- This configuration is only active, if this XML file is placed in the mentioned folder AND if **<isActive>** is set to true. Otherwise, the behavior is the same as before, means with every dictionary synchronization ALL rule executor objects are disposed and created lazily on demand.
- The **<dictionary>** element defines, which rules (defined by the **<ruleIdentifier>**s within the **<affectedRule>** element) are affected by a specific dictionary (specified by the **<dictionaryIdentifier>**). You can define as many **<dictionary>** elements as you wish.
- The **<dictionaryIdentifier>** has to be copied from the dictionary view (part of the Dictionary perspective) in the PIM Desktop client.
- The **<ruleIdentifier>** is composed of the folders and subfolders of the IDQ rule and the name of the IDQ rule itself. This information can be found through the Data quality perspective in the PIM Desktop Client:
- When a dictionary, which is configured here, is synchronized, ONLY the executor objects of the configured affected rules are destroyed and recreated lazily on demand - all other rule executor objects are not hit.
- When a dictionary, which is NOT configured here, is synchronized, the behavior is the same as before, means ALL rule executor objects are disposed and created lazily on demand.



Be aware, that as soon as a specific dictionary is configured in the **DictionaryRuleConfiguration.xml** file, synchronization of that dictionary NEVER leads to a destruction of not affected rule executor objects. That means that if those "not affected" rules executor objects might anyhow have dependencies to the configured dictionary, they might work on a very old state of this dictionary, which might lead to an inconsistent behavior when executing IDQ rules.

i The **DictionaryRuleConfiguration.xml** configuration is only loaded once during server startup. Means, after changing that file, the PIM server has to be restarted, in order to activate the changes.

If this something goes wrong, during loading of this configuration file (e.g. the corresponding **DictionaryRuleConfiguration.xsd** is not existing in the /configuration/HPM/dataquality folder, or there is an syntax error in the **DictionaryRuleConfiguration.xml** file), an ERROR is logged in the server's log file. But still the server starts up and the behavior is the same as before, means with every dictionary synchronization ALL rule executor objects are disposed and created lazily on demand.

2.10 Check integrity of attribute names in key language

The name of an attribute must exist in the key language otherwise errors may occur during operation. The key language can be specified in the server.properties.

If the key language is changed, it is possible that attributes exist that do not have a name in the key language. The following database script can be used to check if all attributes have a name in the key language.

2.10.1 MSSQL

The script for MASTER and SUPPLIER

```
SELECT COUNT(aa1.ID)
FROM [ArticleAttribute] aa1
INNER JOIN [ArticleAttributeLang] aal1 ON aa1.ID = aal1.[ArticleAttributeID]
WHERE aa1.ID NOT IN (
    SELECT aa2.ID
    FROM [ArticleAttribute] aa2
    INNER JOIN [ArticleAttributeLang] aal2 ON aa2.ID = aal2.[ArticleAttributeID]
    WHERE aal2.[LanguageID] = :keyLanguageID )
```

The script for MAIN

```
SELECT COUNT(sga1.ID)
FROM [StructureGroupAttribute] sga1
INNER JOIN [StructureGroupAttributeLang] sgal1 ON sga1.ID = sgal1.
[StructureGroupAttributeID]
WHERE sga1.ID NOT IN (
    SELECT sga2.ID
    FROM [StructureAttribute] sga2
    INNER JOIN [StructureGroupAttributeLang] sgal2 ON sga2.ID = sgal2.
[StructureGroupAttributeID]
    WHERE sgal2.[LanguageID] = 9 )
```

2.10.2 Oracle

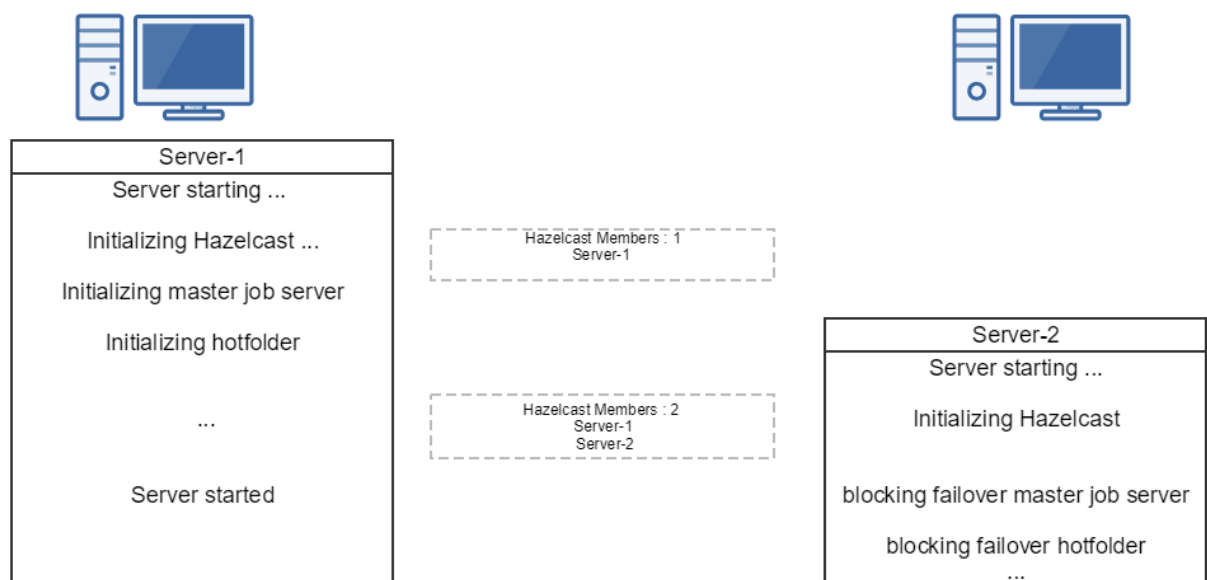
The script for MASTER and SUPPLIER

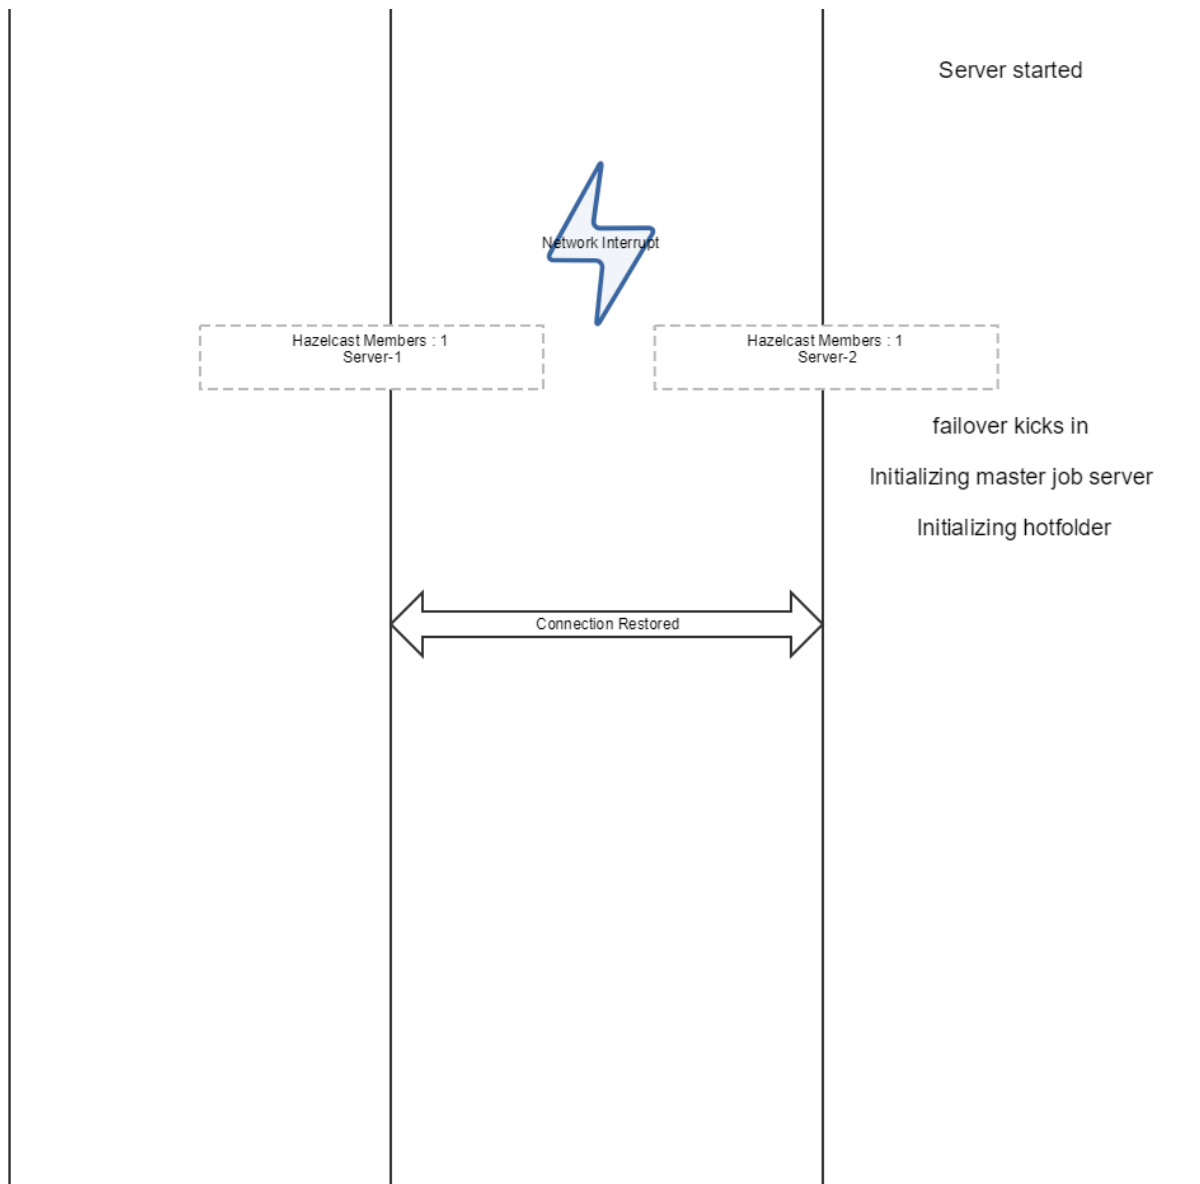
```
SELECT COUNT(aa1.ID)
FROM "ArticleAttribute" aa1
INNER JOIN "ArticleAttributeLang" aal1 ON aa1.ID = aal1."ArticleAttributeID"
WHERE aa1.ID NOT IN (
  SELECT aa2.ID
  FROM "ArticleAttribute" aa2
  INNER JOIN "ArticleAttributeLang" aal2 ON aa2.ID = aal2."ArticleAttributeID"
  WHERE aal2."LanguageID" = :keyLanguageID )
```

The script for MAIN

```
SELECT COUNT(sga1.ID)
FROM "StructureGroupAttribute" sga1
INNER JOIN "StructureGroupAttributeLang" sga1l ON sga1.ID =
sga1l."StructureGroupAttributeID"
WHERE sga1.ID NOT IN (
  SELECT sga2.ID
  FROM "StructureAttribute" sga2
  INNER JOIN "StructureGroupAttributeLang" sga2l ON sga2.ID =
sga2l."StructureGroupAttributeID"
  WHERE sga2l."LanguageID" = 9 )
```

2.11 Network Recovery





2.12 Monitoring with Micrometer

2.12.1 General monitoring approach

The Product 360 server exposes several metrics leveraging the micrometer library (<https://micrometer.io/>). These monitoring systems are supported by the Product 360 server:

- CloudWatch
- Graphite
- NewRelic
- Dynatrace
- Elastic

- Prometheus

Configure the monitoring system you want to use in the micrometer.properties file located in **<PIM_SERVER_INSTALLATION_ROOT>\server\configuration\HPM\micrometer.properties** of the server installation package to expose metrics. The application server needs to be restarted in order to have changes take effect. A list of JVM default metrics will be exposed as well as Product 360 custom metrics. Find a list of the available custom metrics below.


2.12.2 List of Product 360 custom metrics

Meter name in prometheus endpoint	Tags	Description
queue_messages_read_total	queueName - Name of the monitored queue	Number of messages this Product 360 server has read from the queue with queueName
queue_messages_read_success_total	queueName - Name of the monitored queue	Number of messages this Product 360 server has successfully processed after it read them from the queue with queueName
queue_messages_read_error_total	queueName - Name of the monitored queue	Number of messages this Product 360 server has read from the queue with queueName, but errors occurred when they were processed
queue_messages_write_total	queueName - Name of the monitored queue	Number of messages this Product 360 server has written to the queue with queueName
queue_messages_write_error_total	queueName - Name of the monitored queue	Number of messages this Product 360 server has tried to write to the queue with queueName, but failed to do so

2.12.3

Sample for monitoring a Product 360 server with Micrometer, Prometheus and Grafana

Architecture overview to monitor Product 360 servers with Micrometer, Prometheus and Grafana.



The following steps are required to enable monitoring of Product 360 server with Micrometer, Prometheus and Grafana.

- Enable the Prometheus endpoint for the Product 360 server you want to monitor by adjusting the values in the file `micrometer.properties`:

✓ Micrometer settings in micrometer.properties

```
#####
####
### Prometheus settings
###
#####
####
prometheus.enabled           = true
prometheus.uriPath           = /prometheus
prometheus.port              = 9090
prometheus.descriptions      = true
#prometheus.step              =
```

- As a minimum enable prometheus, set the uriPath and the port number
- Restart the Product 360 server
- Verify metrics are available by accessing the Prometheus endpoint you specified. In this sample `http://<your_server_hostname>:9090/prometheus`
- Use a Prometheus server to pull the metrics from the Product 360 endpoints and to write them into a database. You need to add the Product 360 Prometheus endpoint as a scrape target to the `prometheus.yml` file. For details check the Prometheus configuration documentation
- Use a Grafana server to visualize the collected metrics data in Prometheus. Therefore Prometheus is added as data source to Grafana and the meters are queried using PromQL (Prometheus Query Language). To visualize the JVM default metrics you can import one of the available Grafana template dashboards as for instance: `https://grafana.com/grafana/dashboards/4701` To visualize the Product 360 custom metrics you can build your own dashboard and query them by the names listed in the custom metrics table or you can import the Product 360 dashboard containing the Product 360 custom metrics.
Product_360_Grafana_Dashboard.json Specify your Prometheus server as data source during import.

3 Desktop Operation

Combines maintenance, tuning and troubleshooting for the PIM - Desktop

- [Maintenance Tasks](#)(see page 90)
- [Tuning](#)(see page 90)
 - [Performance optimization of the "Documents" view](#)(see page 90)
 - [Filter](#)(see page 91)
 - [Import](#)(see page 92)
- [Trouble Shooting](#)(see page 93)
 - [User Interface is not responding any more \(client freeze\)](#)(see page 93)
 - [Find the problem on the client](#)(see page 93)
 - [Find the problem on the server](#)(see page 93)
 - [Connection timeouts between Server and Client](#)(see page 94)
 - [Rich Text Editor](#)(see page 94)
 - [Unable to use the clipboard actions "Cut", "Copy" and "Paste"](#)(see page 94)
 - [Workaround to prevent the "Invalid Certificate" error \(Mozilla\)](#)(see page 95)
- [Product Paradim \(EGD\) limitations](#)(see page 96)

3.1 Maintenance Tasks

- Backup strategy
- Workspace configuration

3.2 Tuning

3.2.1 Performance optimization of the "Documents" view

PIM - Server shows in its Documents view the content of a media asset category delivered by the integrated media asset provider.

However, the respective media asset providers vary in the data they deliver for a media asset and in the reaction time/speed.

For instance, it might be a problem to access a preview image at provider "A", or the calculation of an image dimension at provider "B" is a problem.

Therefore, the following requirement to media asset providers has been introduced:

- The media asset provider must deliver for each media asset an (unique) identifier, the name, and size in high speed.

If this requirement is not fulfilled, it makes no sense to integrate such a media asset provider.

Furthermore, the load of this view has been optimized so that it is based on the properties returned by the respective media asset provider:

When you are opening a media asset category, a list of all media assets contained by this category is requested from the media asset provider.

The media asset provider has to return the identifiers, the names, and the file sizes immediately.

The values for all other columns the media asset provider did not immediately return can be returned with a "delayed" status.

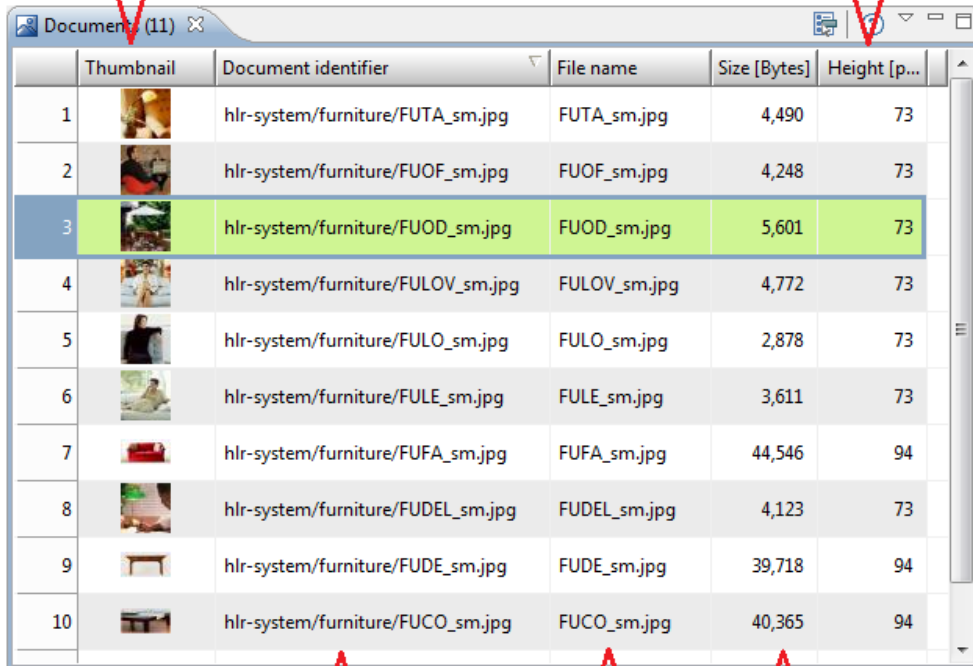
In this case, the view requests these values for each row separately, when the view "scrolls" in the visible field area.











The performance of the view can be improved by hiding unimportant columns.

It is also possible to hide the preview image column. This increases the performance even more.


asynchronous loading

asynchronous loading, if delayed by the media asset system



	Thumbnail	Document identifier	File name	Size [Bytes]	Height [p...
1		hlr-system/furniture/FUTA_sm.jpg	FUTA_sm.jpg	4,490	73
2		hlr-system/furniture/FUOF_sm.jpg	FUOF_sm.jpg	4,248	73
3		hlr-system/furniture/FUOD_sm.jpg	FUOD_sm.jpg	5,601	73
4		hlr-system/furniture/FULOV_sm.jpg	FULOV_sm.jpg	4,772	73
5		hlr-system/furniture/FULO_sm.jpg	FULO_sm.jpg	2,878	73
6		hlr-system/furniture/FULE_sm.jpg	FULE_sm.jpg	3,611	73
7		hlr-system/furniture/FUFA_sm.jpg	FUFA_sm.jpg	44,546	94
8		hlr-system/furniture/FUDEL_sm.jpg	FUDEL_sm.jpg	4,123	73
9		hlr-system/furniture/FUDE_sm.jpg	FUDE_sm.jpg	39,718	94
10		hlr-system/furniture/FUCO_sm.jpg	FUCO_sm.jpg	40,365	94

synchronous loading

 Note: Thumbnails of pdf documents are not displayed in the documents view!

3.2.1.1 Filter

The Documents view provides an interface enabling to filter the content on e.g. JPEG documents. To do this, you have to register an implementation of the Interface

```
com.heiler.ppm.mediaasset.ui.api.IFileDataFilter
```

by implementing the method

```
com.heiler.ppm.mediaasset.ui.views.DocumentsOfCategoryTableView.registerFileDataFilter(IFileDataFilter)
```


of the view. This filter will then be called every time the view is populated with data. From a technical point of view this filter will be called inside the respective IContentProvider implementation.

```
public interface IFileDataFilter
{
    public FileData[] filter( FileData[] fieldTofilter );
}
```

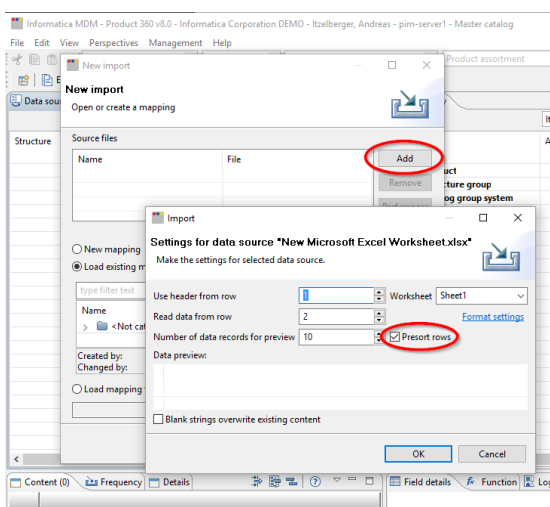
3.2.2 Import

When adding CSV or Excel files to an import project, per default all rows are read into heap memory in order to presort them. This speeds up processing if the data for one item is stored in rows that are spread. So when adding large CSV or Excel data files, consider to increase the client's heap size accordingly.


If presorting is not needed (because the data files are already presorted) or desired, it can be disabled by opening the "Advanced settings" dialog after selecting the CSV file, and here unchecking the checkbox "Presort rows". See also the following screenshot:



For Excel data files find the "Presort rows" checkbox directly in the Settings dialog, after selecting the Excel file:



This "Presort rows" settings only affects the adding of data files to the import project in P360 Desktop Client, but not the import itself. For disabling presorting rows in the import itself, please see the chapter "[Server Operation](#)(see page 45)" (of the Installation and Operation guide) and here the sub chapter "Import", "Preferences regarding memory usage".

 Adding large data files to an import project might become slower, when presorting is deactivated (see setting above) and the import data file is not already presorted.

3.3 Trouble Shooting

- Log Files
- Workspace log
- Server connection lost/reconnect

3.3.1 User Interface is not responding any more (client freeze)

The currently supported operating systems of the Heiler Product Manager do only support a single GUI thread. It is responsible for reading and dispatching operating system events like mouse clicks and keyboard strokes. Usually the main thread reads an event from the operating system, and dispatches this event further to whoever listens on it (in our case SWT widgets). It then waits for all those event listeners to complete and then read again the next events from the operating system (the read-and-dispatch loop is synchronously). It performs this in an endless loop as long as the application is running.

In the PIM - Desktop this thread is also the `main` thread, which makes it quite easy to identify. Therefore, whenever the client stops responding it must be related to the client's `main` thread being busy with other things then reading and dispatching OS events. For example it's waiting on some operation to complete.

3.3.1.1 Find the problem on the client

To pinpoint the reason for the freeze you will need to know what the client is doing (or better waiting for) at the moment of the freeze. For this you will need a [How to create a thread dump](#)(see page 124) of the client.

1. Perform a thread dump of the client at the moment of the freeze
2. Find the `main` thread in the dump and check the stack trace - see where it ends.
3. In case the `main` thread waits on the completion of a server request you will need to move your investigation to the server.
In this case the stack trace of the client's `main` thread usually stops somewhere in the communication framework, e.g. `ClientNodeImpl(AbstractNode).sendRequest(NodeIdentifier, Message, long)`
4. In case the stack trace makes absolutely no sense for you, don't panic, that's not unusual 😊. With the information you collected we should be able to help you identify and fix the problem.
5. In case the client's `main` thread is waiting for a server request to complete, you can move on step further and check the server's threads.

3.3.1.2 Find the problem on the server

Since the server has no user interface, the main thread of the server is more or less irrelevant for this task. All requests from the communication framework are handled by the communication worker threads. The only

problem is to find the right one. Currently you can't determine this by just looking at the thread names, you really have to check all communication framework threads. Communication worker threads are named `CommunicationWorker-<Number>`. Most of them should be in parking position, waiting for a request of a client (e.G. `Unsafe.park(boolean, long)` line: not available [native method]), those can be ignored, of course.

i If you take multiple thread dumps of the server, you can check which thread "didn't move" between those dumps, making it easier to find the responsible one.

When you finally found the communication worker thread which is responsible, it might happen that this thread is itself waiting. Waiting for some other thread or the database. From there on it gets complicated and leads to debugging multi-threading issues and database locking problems. It would go beyond the scope of this wiki, so we would encourage you to contact Heiler Software Support if you need any further assistance on this.

3.3.1.3 Connection timeouts between Server and Client

In case there are connection issues between the server and the client regarding the heartbeat timeout it is possible to change the settings in the `plugin_customization.ini` of the client in the section "Communication CORE Settings".

If these settings do not generate the expected outcome it is also possible to change the rate in which the server checks for messages to the client. As the behavior is that if the server has nothing to send to the client in a specific time frame it will send a heartbeat object to the client. This time frame can be edited in the server's `plugin_customization.ini` if you add the following key:

```
com.heiler.ppm.communication.core/serverPollRate = 10
```

The value is the time in seconds which the server waits for a new message to send it to the client. Default value is 10.

When editing this value please keep in mind that it cannot be lower than the value of `heartbeatInterval`. Otherwise the heartbeat will fail as soon as no activity is send to the client.

If the client do not get bytes in a specific time frame from the server it marks the connection as unusable and the client disconnects from the server.

To disable the heartbeat from client to server in general you can set the `heartbeatInterval` to 0 in the client `plugin_customization` and set the `serverPollRate` in the server `plugin_customization` to 0.

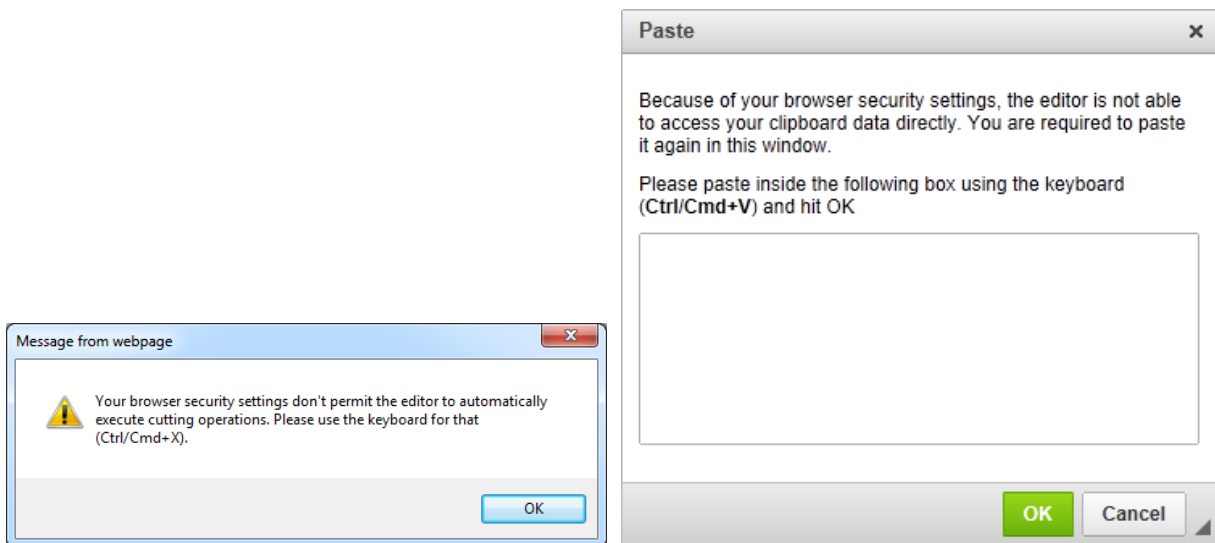
3.3.2 Rich Text Editor

3.3.2.1 Unable to use the clipboard actions "Cut", "Copy" and "Paste"

Due to the browser setting it may be impossible to use the following RichText-Editor buttons



In this case one of the following dialogs may be shown:

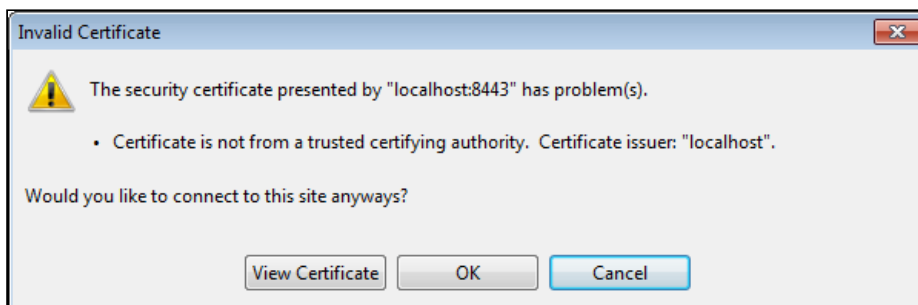


To change this behavior open the "Internet Explorer" and do following steps:

1. Go to Tools -> Internet Options.
2. Go to "Security" tab.
3. Select "Internet" zone, then click on "Custom level..." button.
4. Scroll down to "Scripting" section (at the bottom few).
5. Under "Allow Programmatic clipboard access" option, check or select (tick) to enable

3.3.3 Workaround to prevent the "Invalid Certificate" error (Mozilla)


If the browser control loads any URL which uses an **unsigned** (or self-signed) SSL certificate, an error dialog appears:



Since there is no way to add an "exception" to the XULRunner, the following workaround can be used:

1. Open Firefox and load the domain which uses the invalid certificate (i.e. "https://localhost:8443/")

- The page "This connection is untrusted" will be shown (see below):



- Click on the button "Add exception..." and confirm the dialog which appears
- Go to the "Firefox" profile folder (typically "C:\Users\[USER_NAME]\AppData\Roaming\Mozilla\Firefox\Profiles\[PROFILE_NAME]\")
- Copy the file "**cert_override.txt**"
- Go to the XULRunner profile folder (typically "C:\Users\[USER_NAME]\AppData\Roaming\Mozilla\eclipse\")
- Paste the file "**cert_override.txt**"

Alternatively it is possible to use another workaround:

- Open the file "prefs.js" in the folder "C:\Users\[USER_NAME]\AppData\Roaming\Mozilla\eclipse\".
- Add the following line to the file and save the file:
`user_pref("network.automatic-ntlm-auth.trusted-uris", [host url of IdP server]);`

Next time you start PIM Desktop the browser control will load the URL without any errors.

Alternatively just use a "trusted" certificate (see https://en.wikipedia.org/wiki/Certificate_authority) or switch the setting `http.client.proxy` in "server.properties" to an HTTP URL (instead of HTTPS).

Startig with 8.0.03 it is possible to login to the desktop client via SAML. In this case the same rules apply to the certificate of the SAML IDP server. Generally it is recommended to use a trusted certificate, however If necessary the workaround with the `cert_override.txt` can be applied here as well.

3.4 Product Paradim (EGD) limitations

Please note that the EGD (Extended Generic Data) based product paradigm, which was used in the past for the first PIM Product paradigm, ist now with PIM 7 no longer in Support. The EGD itself will remain as basis for custom entities, but without the `EGDAttributeMapping` entity.

This information is also about some sideeffects in the PIM 7 user interface, yet. For example you are having in the organization perspective double menu entries for product assortments. One is for the old, egd based, product assortment, one for the new product assortments. Unfortunately you cannot see the difference in one glance. Workaround until this is removed with PIM 8 is to try out default object rights on your product assortment. If it does not have the adjusted effect you have to choose the other selection in the default object rights combo box.

4 Control Center Manual

- [Prerequisite\(see page 97\)](#)

- [Function](#)(see page 97)
- [Login](#)(see page 98)
- [Dashboard](#)(see page 98)
 - [Widget creation using templates](#)(see page 98)
 - [Customize widgets](#)(see page 99)
- [Monitoring](#)(see page 100)
 - [Application Server Node](#)(see page 101)
 - [Server tab](#)(see page 101)
 - [Server Nodes](#)(see page 102)
 - [Server Node Details](#)(see page 102)
 - [Clients](#)(see page 103)
 - [Log](#)(see page 104)
 - [Thread Dump](#)(see page 104)
 - [SNMP](#)(see page 105)
- [Deployment](#)(see page 105)
 - [Installation](#)(see page 106)
 - [License](#)(see page 107)
 - [Configuration](#)(see page 107)
- [Support](#)(see page 107)
 - [CSM](#)(see page 108)
- [Additional Features](#)(see page 108)
 - [Opening multiple tabs](#)(see page 108)
 - [Telemetry Actions](#)(see page 109)

4.1 Prerequisite

In order to get the Control Center running you need to follow the installation manual for the Server Installation

4.2 Function

The Control Center is a central application to help to administrate the systems of your cluster. It provides basic functionality like:

- Install and update your Product 360 system in the cluster
- Start server(s)
- Stop servers(s)
- See logs
- See running jobs
- See information of servers
- Gather configuration settings of the system
- Shutdown clients
- Send message to clients
- create thread dumps

4.3 Login

To open the web interface of the Product 360 Control Center you have to know where the service is running. The URL consists of the host name and the port which is defined in the `ClusterixConfig.xml`. By default the port is 9000. So when the service is running on the instance called `myComputer` then you can open the web interface by calling `myComputer:9000`

The login credentials are also defined in the file `ClusterixConfig.xml`. By default you can login with `user:clusterix password:clusterix`.


4.4 Dashboard

4.4.1 Widget creation using templates



Add a Panel...

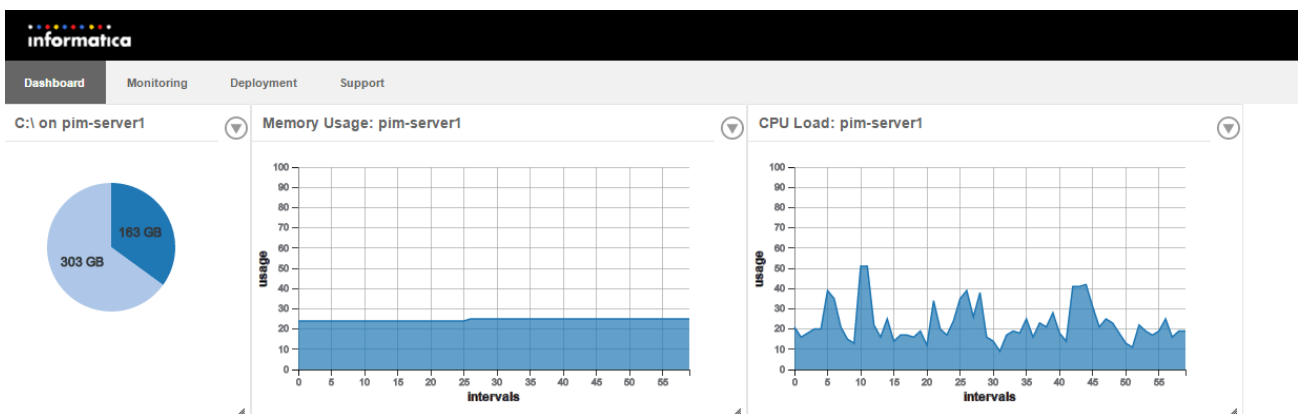
The dashboard is your personal area, where you can pint and save graphical widgets. It allows you to monitore data of a selected server. There are a number of widgets which may help you doing so. If you enter your dashboard the first time, you will be greeted with an empty panel. To add widgets to your dashboard, you can either select the label in the middel of the screen or left click on the **dashboard** tab in the top left corner.



Doing so opens up a window, giving you the opportunity to add your desired panel. Selecting your wanted panel template and left clicking on the 'ok' button, places the selected widget on your dashboard.

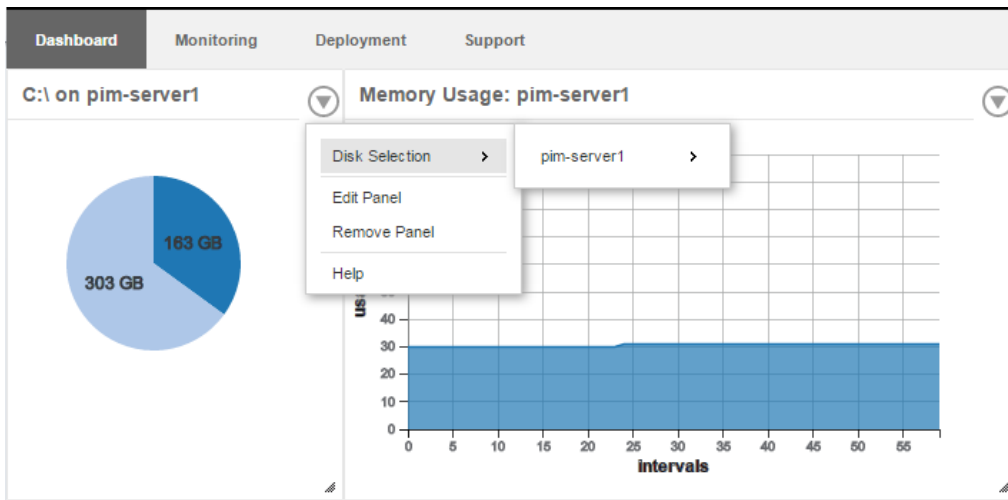
4.4.2 Customize widgets

The dashboard saves your current state and allows you to view your data at any given moment.



To customize your personal dashboard, you can select the arrow symbol in the top right corner of every single widget. Doing so allows you to:

- **select a server** that you want to monitor
- **edit the panel** itself
- **remove** the wanted widget.



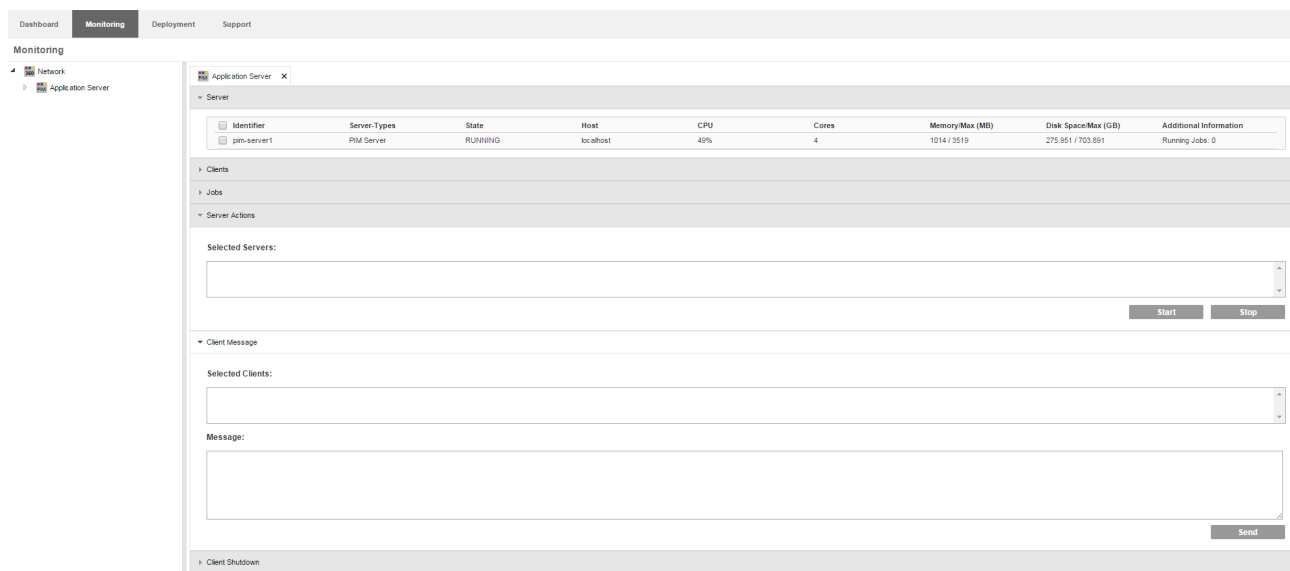
Every change will be saved once you leave your dashboard and is loaded up when you enter it, making it your own personal dashboard.

Note: The Memory Usage widget displays the total Memory **available to the Java Runtime Environment** and the currently allocated Memory available to current and future objects inside this environment at the given point in time of the selected Server **Application**. It does not show how much memory is available to the selected Server Machine itself.

4.5 Monitoring

The Monitoring tab allows you to monitor all systems inside your network. From here on out you are able to watch different categories of all application servers that are currently available in your network system. The left side of the monitoring tab contains a navigation tree, which is displaying the network itself and all available features for each server as nodes. The right side contains all associated data of the selected node inside the navigation tree.

4.5.1 Application Server Node



Clicking on the **Application Server** tab, opens the application server view. This tab contains information like:

- All running application servers and their status
- All connected clients
- All running Jobs that are currently running on the servers

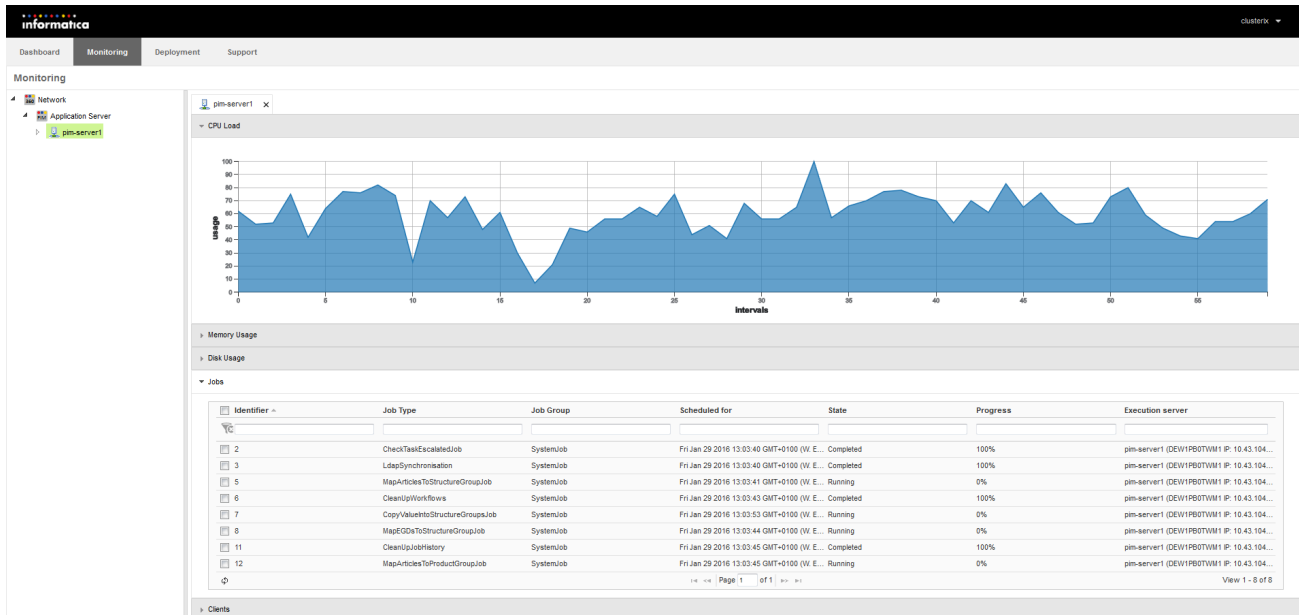
You are also able to select a running server and **start** or **stop** it. If there are clients connected to a server, you can select a client from the **Client Message** tab and send them a message. In the **Client Shutdown** tab you are able to select a client and shut it down.

4.5.1.1 Server tab

The Server tab presents server data about every configured server inside the current network. Together with the general data, which can be seen in the graphic above, you can view the current memory consumption of each server. The shown memory consumption describes how much memory has been allocated to the PIM server **application** running on the system and how much that application is currently using. The allocated memory on each server application can increase automatically if the application finds it necessary. This depends on how much free physical memory is available on the system and if the application is in need of more memory.

The example above has a system 'pim-server1' with a certain amount of physical memory. Currently about 3.5 GB memory has been allocated to the application. The application itself is using about 1.7 GB memory of the allocated 3.5GB memory. An external system monitor like Windows Task manager, may show a memory usage of 3.5GB as this is the memory reserved for the application.

4.5.2 Server Nodes



By opening the application server list and left clicking on one of your running servers inside the navigation tree, you open up a detailed view of the selected server. From here on out you have a wide variety of options and can view the following data:

- A statistical view of the **CPU load** of the currently selected server
- A statistical view of the **Memory Usage** of the currently selected server
- A statistical view of the **Disk Usage** of the currently selected server
- A list of all currently **running jobs** and their status of the selected server
- A list of **connected clients** to the currently selected server

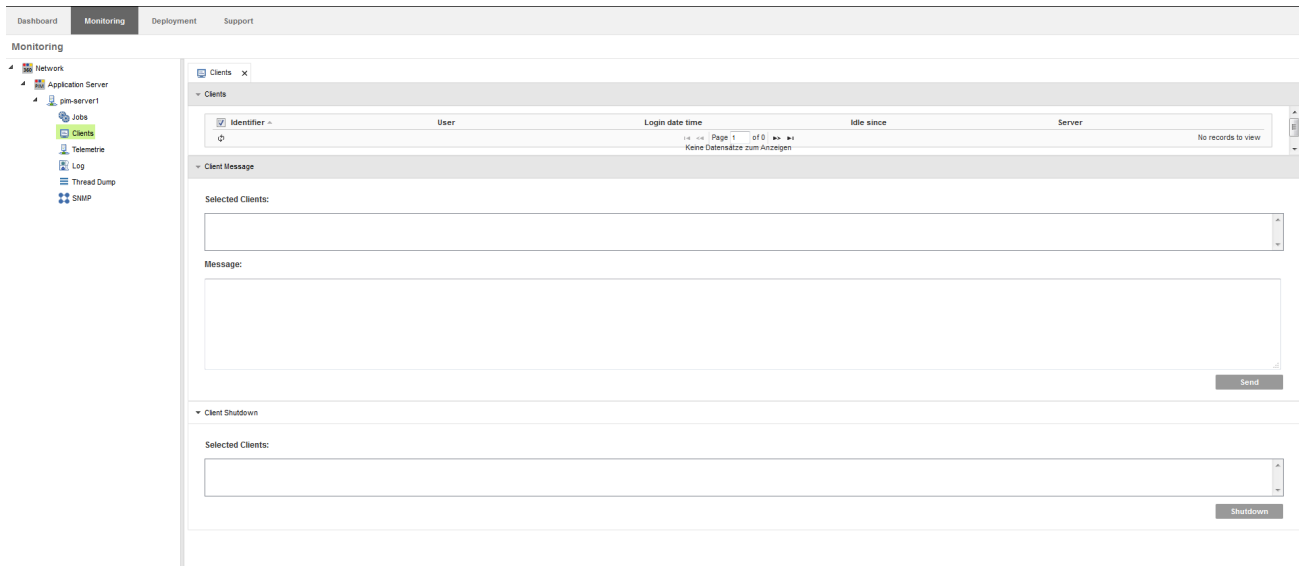
Just like before, you can start and stop the selected server inside the **Server Actions** tab.

Note: The Memory Usage Tab shows how much Memory has been given to the associated Application of a Server and how much is currently in use by the Application itself.

4.5.3 Server Node Details

Extending the selected server node, gives you the option to have a detailed look at the servers actions and data.

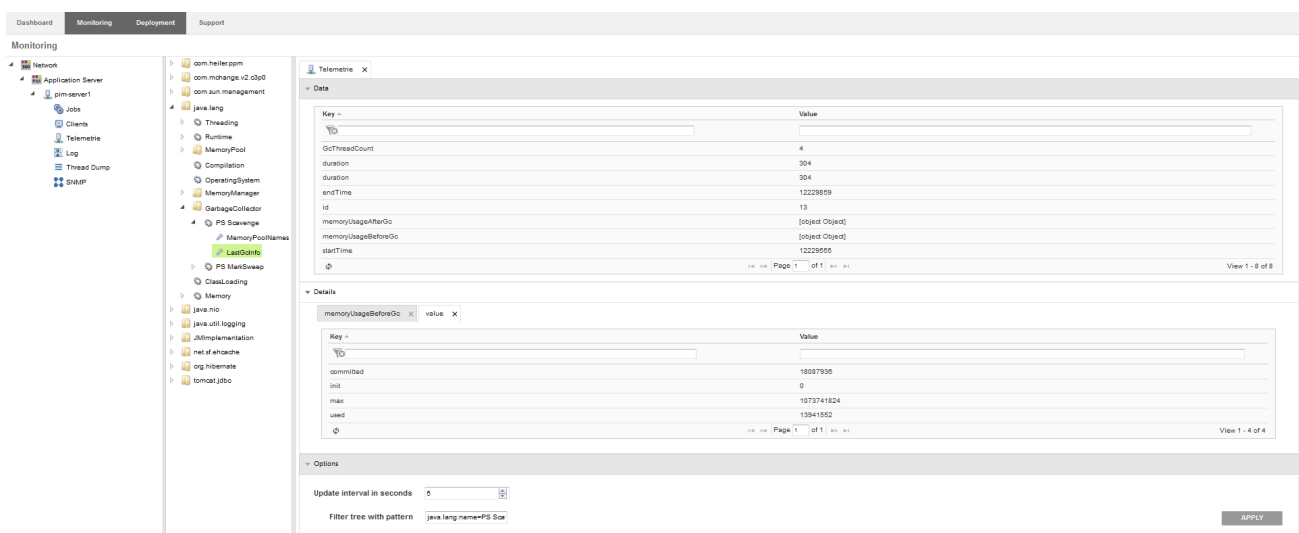
4.5.3.1 Clients



The **Clients** node shows you once again, which clients are connected to the selected server. From here on out you can write messages to a selected client or shut him down.

i Clients, which are connected via web are also shown, but sending messages or the shutdown of web-clients is not supported.

Telemetry



Inside the **Telemetry** node, you have a wide variety of options. The navigation tree allows you to view all kind of telemetry data. By clicking on a node inside the telemetry navigation tree, you can view dynamically changing data inside the **Data tab**. This data is being updated in certain time intervals. To change these

time intervals, simply change the corresponding textfield inside the actions tab and hit enter or select the **APPLY** button at the bottom right side of the screen.

The pattern field allows you to filter the navigation tree according to your pattern. Simply type in a pattern with the following form: **Domain : AttributeName1 = AttributeValue1, . . . , AttributeNameN = AttributeValueN**. Filter the tree by hitting enter or selecting the **APPLY** button.

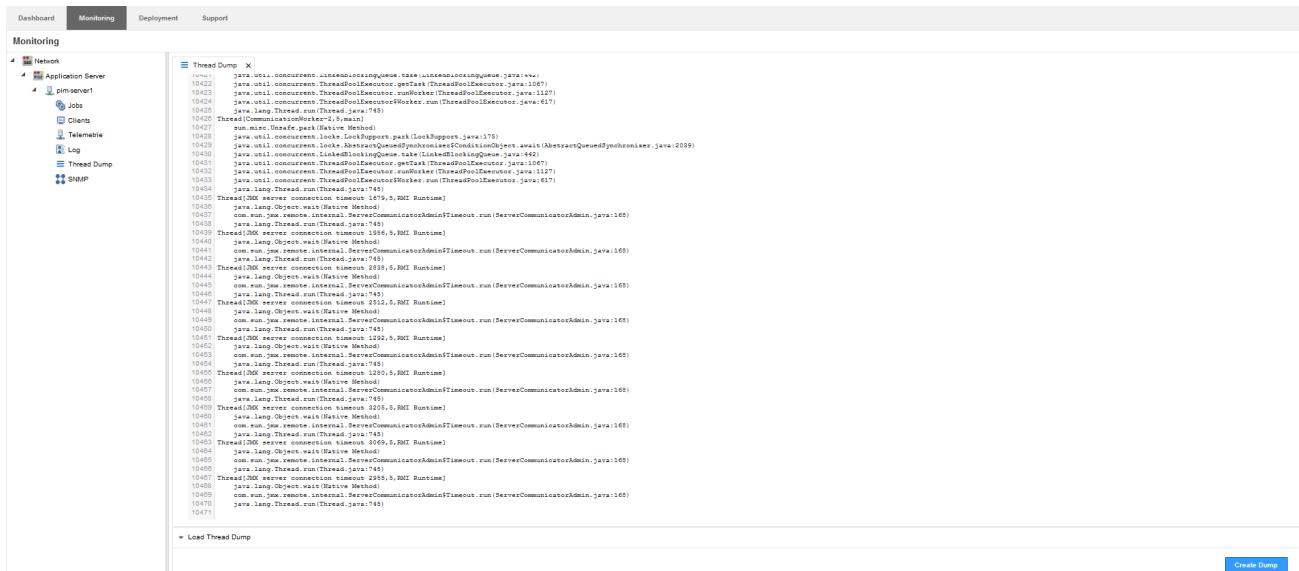
You open up the Detail View by selecting data that is marked as [Object object] inside the **Data tab**. Here you can get more detailed information about the selected entity. If you now select data that is marked as [Object object] inside the **Detail tab**, a new tab opens up and you can view that information inside of it.

4.5.3.2 Log



The **Log** node allows you to view the log of the selected server. While the **Text Content Tab** periodically updates and shows the tail of the log of the selected server, you are able to download the complete log file and the displayed log file by simply selecting the corresponding buttons at the right hand side of the **Log Actions tab**.

4.5.3.3 Thread Dump



Selecting the **Thread Dump** node lets you view the current Thread dump of the selected server. Simply select the **Create Dump** button at the right hand side inside the **Load Thread Dump tab**. The current Thread dump will then be created inside the **Thread Dump tab**.

4.5.3.4 SNMP

The screenshot shows the 'Monitoring' tab in the Informatica MDM interface. On the left, a tree view shows the hierarchy: Network > Application Server > pms-server1 > Jobs > Clients > Telemetry > Lng > Thread Dump > **SNMP**. The main panel displays the 'SNMP' configuration with a table of system metrics.

Name	Object ID	Value
remote command	1.3.6.1.4.1.200.3.0	
description	1.3.6.1.4.1.200.1.1.1	
NonHeapMemoryUsage	1.3.6.1.4.1.200.1.1.2.1	183
HeapMemoryUsage	1.3.6.1.4.1.200.1.1.3.1	1801
SystemCpuLoad	1.3.6.1.4.1.200.1.1.4.1	0.10749979527239539
ProcessCpuTime	1.3.6.1.4.1.200.1.1.5.1	0
SystemLoadAverage	1.3.6.1.4.1.200.1.1.6.1	-1.0
SystemCpuLoad	1.3.6.1.4.1.200.1.1.7.1	0.10749979527239539
TotalCompilationTime	1.3.6.1.4.1.200.1.1.8.1	100437
TotalLoadedClassCount	1.3.6.1.4.1.200.1.1.9.1	19029

Page 1 of 4 | View 1 - 10 of 40

To view data with the **Simple Network Management Protocol**, you need to follow the configuration instructions at the following location : Installation and Operation / Operation / Server Operation / [SNMP Monitoring](#)(see page 57).

Selecting the **SNMP** node, allows you to view all available data of the selected server using SNMP. You can view all available data and its corresponding **Object ID**(OID) which you can later use to monitor your system, using SNMP.

4.6 Deployment

Inside the deployment tab you have the opportunity to configure and administrate deployed systems inside your network. You can:

- administrate and install selected server installation files
- set or change the license file and license key
- edit server configuration files remotely

4.6.1 Installation

The screenshot shows the Informatica MDM - Product 360 10.1- Operation Deployment tab. The 'Installation' node is selected in the sidebar. The main content area displays the 'Installation' sub-tab. The 'Server State' table shows one entry: 'pim-server1' on 'localhost' with 'no state'. The 'Installation Files' table is empty. The 'File Actions (Installation Files)' section includes a 'File upload' area with a 'Drop file here' button and a 'Choose File...' button, and a 'Selected Files' list. The 'Install Actions' section includes a checkbox for 'Use service log on credentials', a 'Username (domain\username):' field, a 'Cluster:' field, a 'Password:' field, and an 'Install' button.

Selecting the Installation node allows you to inspect the installation of the running systems. You can view all running servers at the **Server State tab**. The **Installation Files tab** lets you administrate all currently used installation files. Inside the **File Actions tab** you have the options to chose installation files yourself and upload them. You can also select already used instalation files and either download or delete them, using the corresponding buttons at the bottom right side. The **Install Actions tab** allows you to install PIM using the existing installation files.

There are 2 packages which can be installed by the Control Center a full or a delta package. While the Control Center installs a full package completely, it only exchanges individual files from a delta package. The Control Center uses the filename of the package to distinguish whether it is a full or a delta package. The control center sorts the uploaded packages in a reasonable order. Full packages will be installed first, then the delta packages. If there are multiple full or delta packages uploaded, then the packages will be alphanumeric sorted and installed. By sorting the packages, multiple packages can be uploaded and installed at once

Additionally the Control Center always creates a backup of the current installation and saves the backup into the folder clusterix/install/backups. There is no clean up job which removes old files. This have to be done manually.

For a step by step instruction see the Server Installation documentation. Be aware that the files of the server on which you are connected are used to distribute it to the others.

4.6.2 License

The screenshot shows the 'Deployment' tab selected in the top navigation bar. On the left sidebar, under 'Application Server (Deployment)', the 'License' option is highlighted. The main content area displays the 'License' configuration page. It includes a 'License File' section with a 'Drop file here' text and a 'Choose File...' button, and a 'License Key' section with a 'new License Key' text and an 'Upload' button. A 'Save' button is located at the bottom right.

To change the used license file and license key, the **License** node might be helpful. inside the **License tab** you can simply choose and upload a license file using the corresponding buttons at the right hand side. Pasting a license key inside the license key textfield allows you to use a new license key in your system.

4.6.3 Configuration

The screenshot shows the 'Deployment' tab selected in the top navigation bar. On the left sidebar, under 'Application Server (Deployment)', the 'Configuration' option is highlighted. The main content area displays the 'Configuration' page. It includes a 'Server State' section with a table showing the server's identifier, host, and state. Below this is a 'Configuration Update' section with a checkbox for 'Use service log on credentials', a 'Username (domain\username)' field, a 'Password' field, and an 'Update Configuration' button.

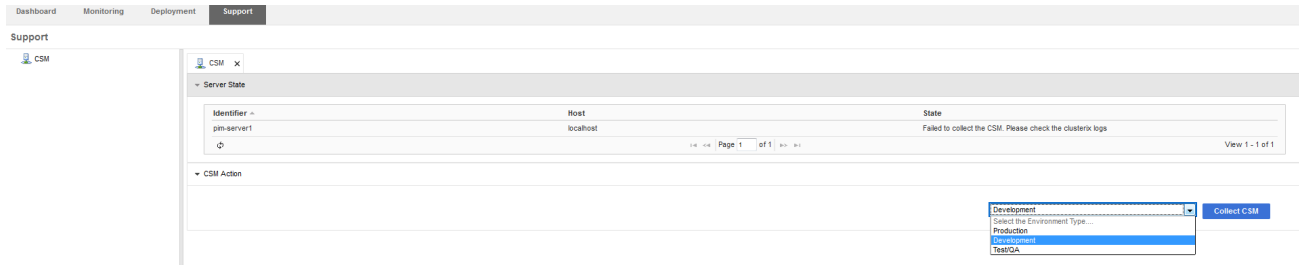
Identifier	Host	State
pim-server1	localhost	no state

The **configuration** node allows you to configure used configuration files of your system. After putting in your credentials inside the **Configuration Update tab**, you can view and configure all kind of system properties inside your network.

4.7 Support

The Support tab features tools to support you collecting all kind of system data.

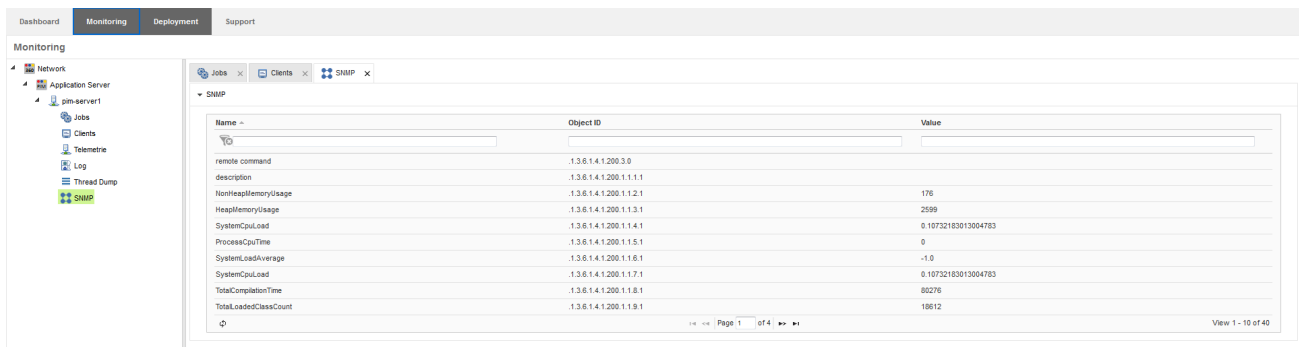
4.7.1 CSM



The CSM node allows you to collect data about running servers in your system. Inside the **CSM Action tab**, select your desired environment type and hit the Collect CSM button. All server will automatically collect data and present it to you.

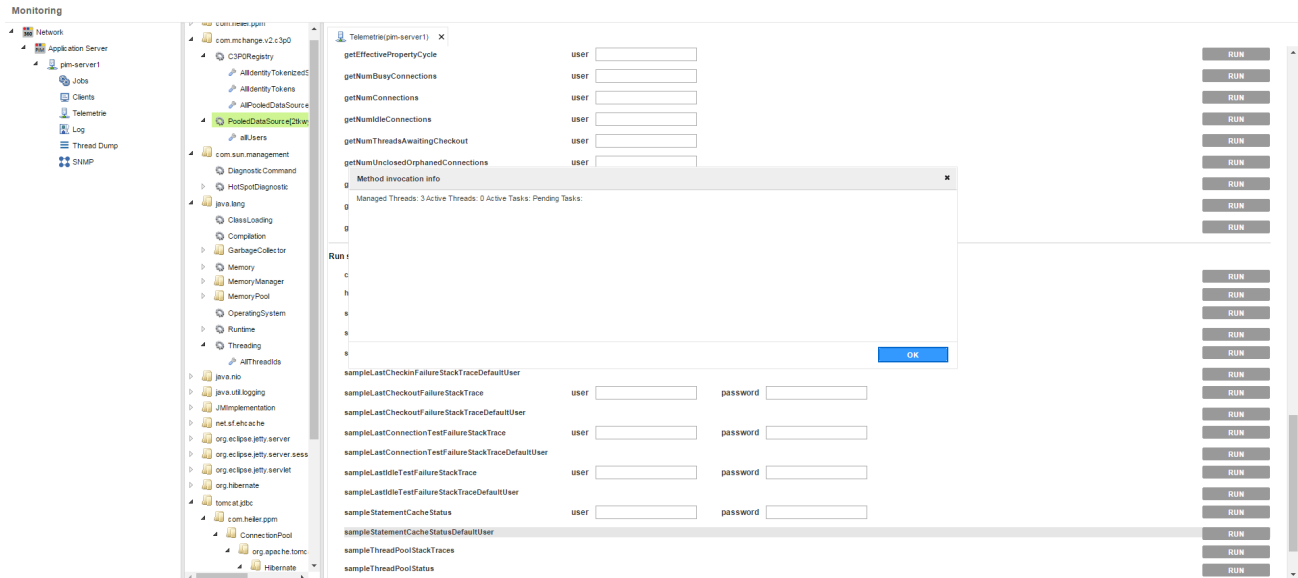
4.8 Additional Features

4.8.1 Opening multiple tabs



Holding Ctrl and selecting a node inside one of the navigation trees, opens up multiple tabs. You can switch between these tabs by simply clicking on them at the top of the screen.

4.8.2 Telemetry Actions



The **Telemetry Actions** tab allows you to change data of the selected Node. By carefully filling out the given fields with valid data, you can change system parameters immediately. There are three different kinds of operations available:

- **Retrieve specific data:** Used to get certain information about an object
- **Change specific data:** Used to change specific object values
- **Run specific data:** Used to run specific operations on the given MBean

5 Audit Trail Operation


This page describes how to backup and restore audit trail data from and to elastic search indexes. It's typically used in environments in which the builtin backup/restore features of elastic can not be used or the audit trail documents should be transferred to a different kind of storage. The exported data is pure JSON and has no dependencies to elasticsearch any more. Scripts and examples of this page work for CentOS Linux environments.

For everyday backup tasks in an on-premises scenario we would recommend to use the out of the box features of elasticsearch. Please refer to the corresponding documentation of Elasticsearch for this.

In a standard operation mode of Product 360, the audit trail processor intercepts every data management request (create/update/delete) and generates an entity item change document JSON which is stored in the Elasticsearch server. If the Elasticsearch server is not accessible, then the transaction is rolled back in Product 360. There might be intermittent network failures between Elasticsearch and Product 360 servers, so there are in-built watchdog threads running in Product 360 on periodic intervals to maintain the eventual data consistency between the product data in Product 360 relational databases and audit trail data in Elasticsearch.

The audit trail storage lifecycle is powered by Elasticsearch's index lifecycle management. Depending on the retention policy defined for an entity in the repository, the lifecycle management is activated.

By default LONG_RETENTION will delete the audit trail data after 2 years and SHORT_RETENTION will delete the audit trail data after 2 months. The retention policies can be configured. Please refer to the configuration manual.

 Customers should take great care in deciding how long they should retain their audit trail data.

5.1 Backup to File

It is always good to archive audit trail data in a human-readable format before it is purged. The following shell script helps in achieving the archival of audit trail data.

Install a third-party library called "Elasticdump"

```

1      * npm (specific to CentOS)
2      1. If npm already installed, goto elasticdump installation
3      2. If npm installed already and want to upgrade to latest version,
      goto step 4
4      3. If npm not installed, goto step 5
5      4. uninstall old version of npm
6         - sudo rm -rf /var/cache/yum
7         - sudo yum remove -y nodejs
8         - sudo rm /etc/yum.repos.d/nodesource*
9         - sudo yum clean all
10     5. install npm version (Example version 15)
11        - curl --silent --location https://rpm.nodesource.com/setup_15.x |
      sudo bash -
12     6. install nodejs
13        - sudo yum -y install nodejs
14
15     * elasticdump
16        - npm install elasticdump@6.31.5 -g

```

audit_trail_backup.sh

```

1  #!/bin/bash
2  # Audit trail backup/archiving script
3  #####
4  # Script Name   : audit_trail_backup.sh
5  #
6  # Description   : Script for archival of elasticsearch audit trail data
7  #                 using elasticdump   #
8  # Args          : entityName, hasOwnIndex, datasource, startDate, endDate
9  #
10 #####
11
12 #####
13 #
14 # Configurable base parameters. Please refer README for usage
15 #
16 #####
17 #
18 declare -r ELASTIC_BASE_URL="http://localhost:9200"
19
20 declare -r BASE_DOWNLOAD_PATH="home/User/AUDITTRAILBACKUP"
21
22 declare -r P360_SYSTEM_NAME="localhost"
23
24 declare -r ELASTICDUMP_BATCH_SIZE=5000
25
26 declare -r ARCHIVE_FILE_SIZE_MBS=100
27
28 declare -r QUIET_MODE=false
29
30 #####
31 #####
32 # Mandatory parameters required to uniquely identify indexes and time
33 # range for archive #
34 # Please do not change them
35 #
36 #####
37 #####
38
39 declare entityName=${entityName}
40
41 declare datasource=${datasource}
42
43 declare hasOwnIndex=${hasOwnIndex}

```

```

38 declare startDate=${startDate}
39
40 declare endDate=${endDate}
41
42
43 while [ $# -gt 0 ]; do
44     if [[ $1 == *"--"* ]]; then
45         param="{1/--/}"
46         declare $param="$2"
47     fi
48     shift
49 done
50
51 #####
52 #####
53 ##
54
55 validate() {
56     if [ -z "$entityName" ] || [ -z "$endDate" ] || [ -z "$startDate" ] ||
57 [ -z "$datasource" ] || [ -z "$hasOwnIndex" ]; then
58         echo "validation failed, Please check mandatory fields and try
59 again!"
60         exit 0
61     fi
62     if [ -z "$(which elasticdump)" ]; then
63         echo "elasticdump not installed, Please refer README file for
64 installation details"
65         exit 0
66     fi
67 }
68
69 getElasticSearchIndexURL() {
70     # elastic indices are in lowercase
71     entityNameLowerCase=${entityName,,}
72
73     if [ $hasOwnIndex == "true" ]; then
74         local elasticSearchIndexName="audit_${entityNameLowerCase}_${
75 {datasource}*"
76         elasticSearchIndexURL="${ELASTIC_BASE_URL}/${P360_SYSTEM_NAME}.${
77 {elasticSearchIndexName}"
78
79     # entities that do not have their own indices are stored in elastic
80 shared indices
81     else
82         local elasticSearchIndexName="audit_*term_${datasource}*"
83         elasticSearchIndexURL="${ELASTIC_BASE_URL}/${P360_SYSTEM_NAME}.${
84 {elasticSearchIndexName}"
85     fi

```



```

82
83     #validate elastic index
84     local httpStatusCode=$(curl -s -w "%{http_code}" -I
"$elasticSearchIndexURL?allow_no_indices=false" -o /dev/null)
85
86     if [ $httpStatusCode != "200" ]; then
87         echo "$elasticSearchIndexURL does not exists"
88         exit 0
89     fi
90
91     # exclude migrate indices from backup
92     elasticSearchIndexURL="$elasticSearchIndexURL,-${P360_SYSTEM_NAME}.${
{elasticSearchIndexName}_migrate"
93 }
94
95 getFolderName() {
96
97     local monthYear=${startDate:0:7}
98     folderName="audit_${entityNameLowerCase}_${datasource:-default}_${
{monthYear}"
99 }
100
101 createFolderIfNotPresent() {
102     mkdir -p "${BASE_DOWNLOAD_PATH}/${folderName}"
103 }
104
105 buildElasticSearchQuery() {
106
107     searchQuery="{\"query\":
108         {\"bool\":
109             {\"filter\" :
110                 [{\"term\" : { \"_entity\" : \"$entityName\" } },
111                 {\"term\" : {\"_relationshipType\":
112                     {\"range\": {\"_eventTimestamp\": {\"gte\": \"$
113 {startDate}\", \"lt\": \"$${endDate}\"}}}}]}]}\"
114 }
115
116 executeElasticdumpCMD(){
117     elasticdump --input=$elasticSearchIndexURL --searchBody="$searchQuery" --
118     output=$BASE_DOWNLOAD_PATH/$folderName/data.json --fileSize=$
119     {ARCHIVE_FILE_SIZE_MBS}mb --sourceOnly --quiet=$QUIET_MODE --
120     limit=$ELASTICDUMP_BATCH_SIZE --fsCompress
121 }
122
123 purgeEmptyDirs(){
124     if [ ! "$(ls -A $BASE_DOWNLOAD_PATH/$folderName)" ]; then
125         echo "$BASE_DOWNLOAD_PATH/$folderName is empty, going to purge"
126         rm -r $BASE_DOWNLOAD_PATH/$folderName
127     fi
128 }
129
130 ##### START #####

```

```

127
128 # validate user input
129 validate
130
131 # construct index complete path
132 getElasticSearchIndexURL
133
134 # download file folder name
135 getFolderName
136
137 # create download folder
138 createFolderIfNotPresent
139
140 buildElasticSearchQuery
141
142 # execute dump command
143 executeElasticdumpCMD
144
145 #delete empty folders
146 purgeEmptyDirs
147
148 ##### END #####

```

How to use audit_trail_backup.sh

Configure below attributes

- ELASTIC_BASE_URL : Elasticsearch url base path e.g. "http://localhost:9200"
- BASE_DOWNLOAD_PATH : directory to keep downloaded data e.g. "/home/Users/audittrailbackup"
- P360_SYSTEM_NAME : system name as defined **in** server properties, Example: "localhost"
- ELASTICDUMP_BATCH_SIZE: size of batch processed **in** one go, recommended value : 5000
- ARCHIVE_FILE_SIZE_MBS: max size of **split** files created **in** mbs e.g. 100
- QUIET_MODE : Set it to **true** to suppress elasticdump logs

This **file** is called from wrapper.sh generally, but it could also be called independently using below **command format**

```

bash audit_trail_backup.sh --entityName "Article" --hasOwnIndex "true" --
datasource "master" --startDate 01-May-2020T00:00:00 --endDate 31-May-2020T23:59:59

```

In order to archive multiple entities, the following shell scripts can be used.

wrapper.sh

```

1  #!/bin/bash
2
3  #####
4  # Script Name   : wrapper.sh
5  #
6  # Description   : wrapper script to call backup script by providing
7  #                 required arguments
8  # Args          : User has to provide entityName, datasource and
9  #                 hasOwnIndex arguments by checking repository settings #
10 # Author        :
11 #
12 # Version       : baseline
13 #
14 #####
15
16 #####
17 #
18 #                 **IMPORTANT**
19 #
20 # User has to provide list of "{entityName}, {datasource}, {ownIndex}" as
21 # it is from repository #
22 # for all entities which require backup
23 #
24 #####
25 ENTRIES=(
26     "Article, master, true"
27     "Article, supplier, true"
28     "Variant, master, true"
29     "Variant, supplier, true"
30     "Product2G, master, true"
31     "Product2G, supplier, true"
32     "StructureGroup, main, false"
33     "Structure, main, false"
34 )
35
36 TRACKER_FILE_BASE_PATH=~/.AUDITTRAILBACKUP
37
38 BACKUP_SCHEDULE_MONTHS=12
39
40 #####
41 #####

```

```

36 # Used to get time range for the data to be backed up. This is calculated
    # by subtracting
37 # BACKUP_SCHEDULE_MONTHS from current month. This is to be specified by
    # the user.
38 # Example - For BACKUP_SCHEDULE_MONTHS = 6 and current month Nov'20
    #
39 # backup start date is 01-May-2020 00:00:00 and end date 31-May-2020
    23:59:59
40 #####
    #####
41
42
43 backupMonthNumber=$(date -d "$(date +%m) - $BACKUP_SCHEDULE_MONTHS month"
    '+%0m')
44
45 backupYear=$(date -d "$(date +%Y-%m-1) - $BACKUP_SCHEDULE_MONTHS month" +%Y)
46
47 backupMonthDays=$(cal $backupMonthNumber $backupYear | xargs echo | awk
    '{print $NF}')
48
49 backupStartDate=$(date -d "$(date +%Y-%m-1) - $BACKUP_SCHEDULE_MONTHS
    month" +%Y-%m-%d'T'00:00:00)
50
51 backupEndDate=$(date -d "$(date +%Y-%m-1) - $BACKUP_SCHEDULE_MONTHS month +
    $(( $backupMonthDays - 1 )) days" +%Y-%m-%d'T'23:59:59)
52
53
54 ##### Creates separate logging folder#####
55 declare executionDateTime="$(date '+%Y-%m-%dT%H:%M:%S')"
56 mkdir -p logs/$executionDateTime
57
58 #####
    #####
59 # Iterating over all entries and calling backup script for each
    entity,datasource,ownIndex
60 # combination.
    #
61 #####
    #####
62
63 for index in "${ENTRIES[@]}; do
64
65     IFS=' ' read -r -a entry <<< "$index"
66     entityName=${entry[0]}
67     datasource=${entry[1]}
68     hasOwnIndex=${entry[2]}
69     (
70         echo "backup script executing for entity: $entityName, datasource:
        $datasource, hasOwnIndex: $hasOwnIndex, startDate: $backupStartDate,
        endDate: $backupEndDate"
71

```

```

72      bash audit_trail_backup.sh --entityName $entityName --hasOwnIndex
      $hasOwnIndex --datasource $datasource --startDate $backupStartDate --
      endDate $backupEndDate
73
74      echo -e $entityName '\t' '\t' $datasource '\t' '\t' $hasOwnIndex '\t'
      '\t' $(date) >> $TRACKER_FILE_BASE_PATH/tracker.log
75      ) 2>&1 | tee -a ./logs/$executionDateTime/audit_${entityName,,}_$
      {datasource}.out
76
77      done

```

How to use wrapper.sh

wrapper.sh - wrapper **for** audit_trail_backup script

Configure below attributes

- ENTRIES : User has to provide list of "{entityName}, {datasource}, {ownIndex}" as it is from repository

for all entities **which** require backup

Example:

```

("Article, master, true"
 "Article, supplier, true"
 "Structure, main, false")

```

- TRACKER_FILE_BASE_PATH : Example: /home/Users/tracker

- BACKUP_SCHEDULE_MONTHS : Specify backup duration e.g. - **for** backup schedule of 6 and current month November'2020 it will archive whole of May'2020

bash wrapper.sh

Periodic backup

The wrapper.sh can be associated with a monthly cron job and every month audit trail data will keep getting archived.

5.2 Restore from File

Once the audit trail data is archived, it can be safely stored elsewhere. It is also easy to restore the audit trail data from the archived format. The following shell scripts help in restoring the archived audit trail data into a different Elasticsearch index.

restore_audit_trail_backup.sh

```

1  #!/bin/bash
2
3  #####
4  # Script Name   : restore_audit_trail_backup.sh
5  #
6  # Description   : restores archived files into elasticsearch
7  #
8  # Args          : archiveFolder path and index name
9  #
10 #####
11
12 ###Configurable base parameters.Please refer README for usage###
13 declare -r ELASTIC_BASE_URL="http://localhost:9200"
14
15 declare -r ELASTICDUMP_BATCH_SIZE=5000
16
17 declare -r INDEX_MAPPING_FILE="./elastic_index_mapping.json"
18
19 declare -r QUIET_MODE=false
20
21 # Mandatory parameters required to loads backup to elastic indices
22 declare archiveFolder=${archiveFolder}
23
24 declare index=${index}
25
26 while [ $# -gt 0 ]; do
27     if [[ $1 == *"--" ]]; then
28         param="${1/--/}"
29         declare $param="$2"
30     fi
31     shift
32 done
33
34 #####
35 # Index creation using json mapping present in INDEX_MAPPING_FILE
36 #
37 #####
38 declare value=$(tr -d '\040\011\012\015' < $INDEX_MAPPING_FILE)
39 curl -XPUT $ELASTIC_BASE_URL/$index -H 'Content-Type: application/json' -d $value
40
41 #####
42 # Iterate over .json extension files under archiveFolder, and
43 #

```

```

39 # loads data from each file to elastic index
40 #
41 #####
42 echo "Starting restoration from $archiveFolder to $index"
43 for filename in $archiveFolder/*.json.gz; do
44     (
45         echo $filename
46         gunzip -c $filename > $archiveFolder/
47         restoration_intermediary.json
48         elasticdump --input=$archiveFolder/restoration_intermediary.json
49         --output=$ELASTIC_BASE_URL --output-index=$index --type=data --
50         transform="doc._source=Object.assign({},doc)" --
51         limit=$ELASTICDUMP_BATCH_SIZE --quiet=$QUIET_MODE
52         rm -rf $archiveFolder/restoration_intermediary.json
53     ) 2>&1 | tee -a log_restore_"$(date +%Y-%m-%d)".out
54 done
55 echo "Restoration complete"
56 ##### END #####

```

Audit Trail Mapping JSON

```

1  {
2    "mappings": {
3      "dynamic": "strict",
4      "properties": {
5        "_changeSummary": {
6          "type": "object",
7          "enabled": false
8        },
9        "_changeType": {
10         "type": "keyword"
11       },
12       "_changedEntities": {
13         "type": "keyword"
14       },
15       "_changedFields": {
16         "type": "keyword"
17       },
18       "_container": {
19         "properties": {
20           "_entityId": {
21             "type": "integer"
22           },
23           "_externalId": {
24             "type": "keyword"
25           },
26           "_internalId": {
27             "type": "keyword"
28           }
29         }
30       },
31       "_entity": {
32         "type": "keyword"
33       },
34       "_entityItem": {
35         "properties": {
36           "_entityId": {
37             "type": "integer"
38           },
39           "_externalId": {
40             "type": "keyword"
41           },
42           "_internalId": {
43             "type": "keyword"
44           }
45         }
46       },
47       "_eventTimestamp": {
48         "type": "date",
49         "format": "strict_date_optional_time_nanos"

```



```

50     },
51     "_identifier": {
52         "type": "keyword"
53     },
54     "_invalidReason": {
55         "type": "keyword"
56     },
57     "_migrationId": {
58         "type": "keyword"
59     },
60     "_module": {
61         "type": "keyword"
62     },
63     "_relationshipType": {
64         "type": "join",
65         "eager_global_ordinals": true,
66         "relations": {
67             "changeSummaryDoc": "triggerFiredDoc"
68         }
69     },
70     "_revision": {
71         "properties": {
72             "_entityId": {
73                 "type": "integer"
74             },
75             "_externalId": {
76                 "type": "keyword"
77             },
78             "_internalId": {
79                 "type": "keyword"
80             }
81         }
82     },
83     "_transactionStatus": {
84         "type": "keyword"
85     },
86     "_triggerStatus": {
87         "type": "keyword"
88     },
89     "_user": {
90         "properties": {
91             "_entityId": {
92                 "type": "integer"
93             },
94             "_externalId": {
95                 "type": "keyword"
96             },
97             "_internalId": {
98                 "type": "keyword"
99             }
100         }
101     }
102 }

```

```
103     }
104     }
```

How to use restore_audit_trail_backup.sh

restore_audit_trail_backup.sh - loads data from backup json files to elastic indices

Configure below attributes

- ELASTIC_BASE_URL : Example: "http://localhost:9200"
- ELASTICDUMP_BATCH_SIZE: size of batch processed **in** one go, recommended value : 5000
- QUIET_MODE : Set it to **true** to suppress elasticdump logs
- DEFAULT_INDEX_MAPPING_FILE : Provide path to default P360 index mapping json **file** e.g **"/home/Users/mapping.json"**

```
bash restore_audit_trail_backup.sh --archiveFolder /home/
audit_product2g_master_2020-10 --index restored_audit_product2g
```

6 Log File Overview

This page gives an overview about the log files and their configuration options.

Module	Default Log File Location	Log File Configuration	Remarks
Application Server Web Client	<ServerInstallationRoot>\server\logs*.log <ServerInstallationRoot>\server\wrapper.log	<ServerInstallationRoot>\server\configuration\HPM\log4j2.xml	Application log files (out, error, perf) Windows Service Wrapper log
Desktop Client	<ClientInstallationRoot>\client\logs*.log	<ClientInstallationRoot>\client\configuration\HPM\log4j2.xml	
Supplier Portal	<SupplierPortalInstallationRoot>\logs <SupplierPortalInstallationRoot>\tomcat\logs	<SupplierPortalInstallationRoot>\configuration\logback.xml	Application Logs Tomcat Logs

Module	Default Log File Location	Log File Configuration	Remarks
Media Manager	<code><MediaManagerInstallationRoot>\Informatica Media Manager\IMM_protocols <MediaManagerInstallationRoot>\OpasGWebServer\ Tomcat\logs <MediaManagerInstallationRoot>\OpasGWebServer \XOBSessionManager\logs</code>		
Web Search	<code><WebSearchInstallationRoot>\log\server.log <WebSearchInstallationRoot>\<tomcat- server>\logs</code>	<code><WebSearchInstallationRoot>\<tomcat- server>\lib\log4j.xml <WebSearchInstallationRoot>\<tomcat- server>\conf\logging .properties</code>	Application Logs Tomcat Logs
Database Setup	<code><DatabaseSetupInstallationRoot>\dbclient\logs\ *.log</code>	<code><DatabaseSetupInstallationRoot>\dbclient\con figuration\log4j2.xml</code>	

6.1 Additional Logging of Dataquality

If you want verbose logging of the execution of a dataquality rule by the rule engine, go to the `log4j2.xml` file and change the priority of the IDQ category like below:

```
<Logger name="IDQ" level="DEBUG" />
```

The log files can also be found in the server log folder `<ServerInstallationRoot>\server\logs\`. Each rule configuration will be executed in it's own worker thread, so a log file will have following name pattern:

```
<JobWorkerName>-<RuleConfigurationName>-<StartDate>.log
```

6.2 Additional Logging of Change Summary and Datagraph

To get the insides of datagraphs and change summaries before store them to the database you can enable the following logger. You can define the logger for each subclass of

```
com.heiler.ppm.std.server.entity.EntityPersistenceManagerImpl
```

This sample enables the logging for `ArticlePersistenceManager`.

```
<Logger  
name="com.heiler.ppm.article.server.internal.entity.ArticlePersistenceManager"  
level="TRACE" />
```

Use this logger wisely. The logging output is big.

6.3 Additional Logging of Media Asset Processing

If you want to fetch the logging of JMS processing, go to the `log4j2.xml` file and add the following setting:

```
<Logger name="PPM.MEDIA_ASSET_NOTIFICATION" level="DEBUG"/>
```

If you want to fetch the logging of Media Manager connector API methods, go to the `log4j2.xml` file and add the following setting:

 Note: Connector logging via PIM is available since PIM 8.0.01

```
<Logger name="com.heiler.imm.connector" level="DEBUG" />
```

For more details information please visit the page [Logging for Media Manager Connector within PIM](#).

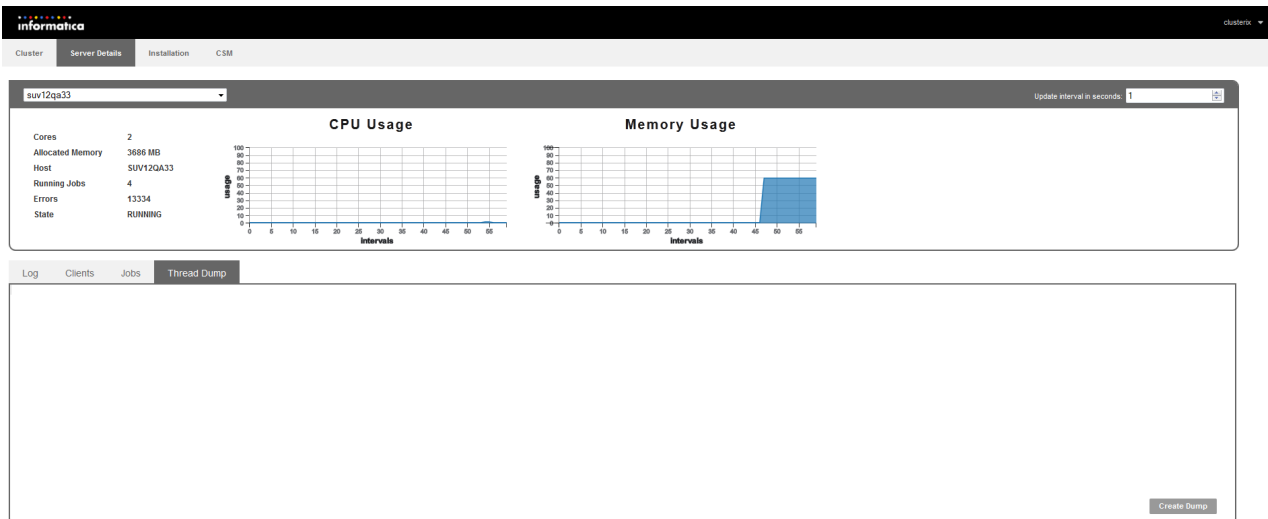
7 How to create a thread dump

This page describes how to obtain a thread dump for the product manager's client and server. A thread dump returns the stack traces of **all threads** of the java virtual machine it's executed on, at the moment of execution.

- [Control Center](#)(see page 125)
- [Console Window](#)(see page 125)
- [VisualVM \(JMX Interface\)](#)(see page 125)
- [Analyzing Thread Dumps](#)(see page 127)

7.1 Control Center

The easiest way to create a thread dump of a running server is by using the Control Center. In the Server details view you can find the Thread Dump tab. By clicking on Create Dump, the dump will be created for the selected server and displayed in the view where it can be analyzed.



7.2 Console Window

When the client or server is being executed in debug mode (using the start_debug.cmd files) you can easily perform a thread dump to the console window. Just press CTRL + BREAK. This will print the stack trace to the console window.

⚠ Starting with PIM - Server version 5.x this also performs a memory dump - which might take a while.

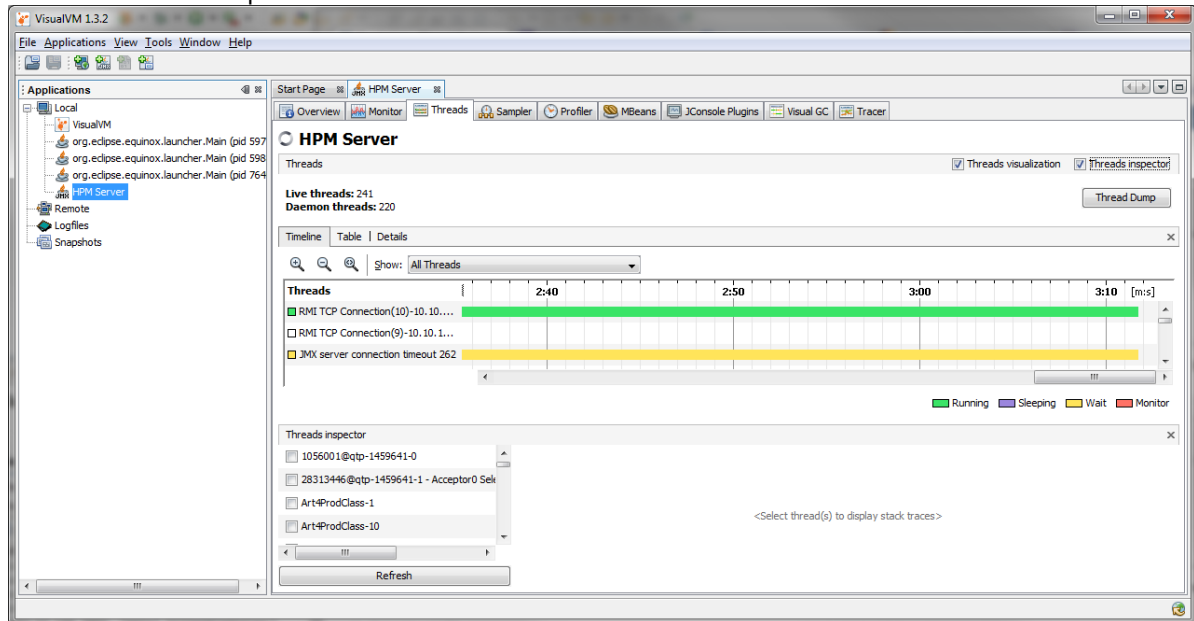
Of course, this is no feasible way for a productive system, which should not be run in debug mode at all!

7.3 VisualVM (JMX Interface)

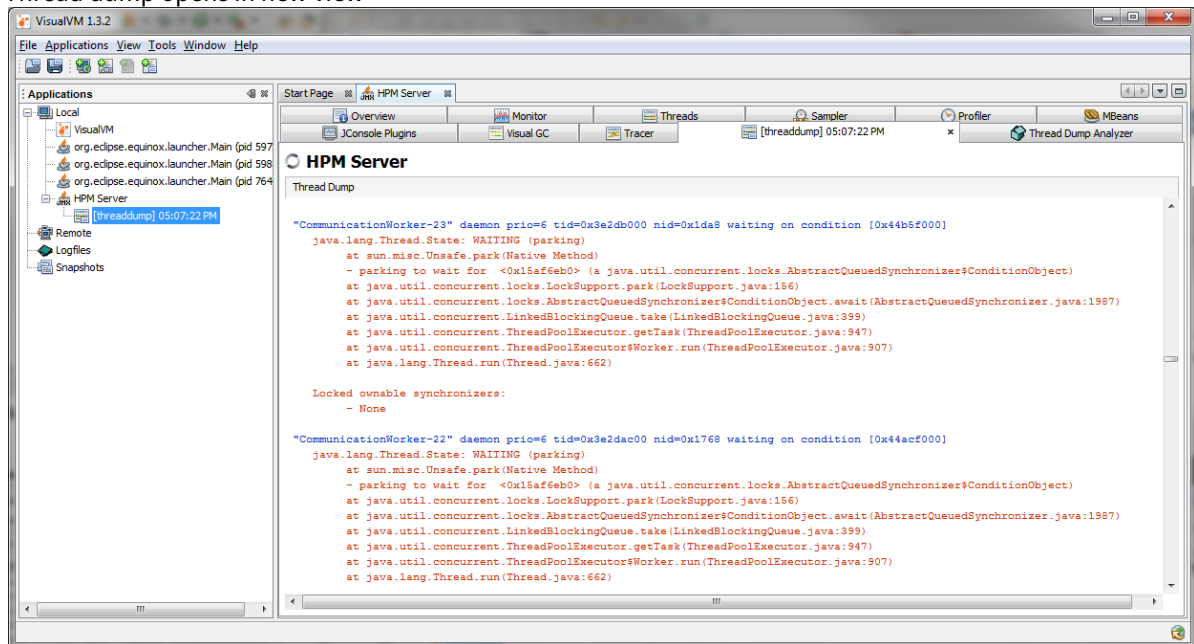
With VisualVM you can connect to a running java virtual machine and perform a thread dump at any given time. Please see the [Tuning Advisory](#)(see page 124) for details on how to enable the JMX (Java Management eXtension) for the server.

1. Open Visual VM
2. [Connect to the server](#)(see page 124)
3. Switch to "Threads" view

4. Click on "Thread Dump"



5. Thread dump opens in new view



⚠ If you have the choice, use Visual VM - it's much easier to work with 😊

If the tool hangs during Thread dump loading, you can also find the written DUMPS in the local user directory at C:\Users\<userId>\AppData\Roaming\VisualVM\8u40\repository\jmxapplications

7.4 Analyzing Thread Dumps

Analyzing Thread Dumps requires some internal knowledge of PIM Server architecture. Here are some useful hints:

- The purpose of thread groups can be usually determined by the thread's name.
 - CommunicationWorker are used by the internal communication framework to process requests from clients
 - http-* are managed by Jetty web server and used for both PIM Web as well as REST calls
- A useful tool for analyzing ThreadDumps is ThreadDumpAnalyzer <https://github.com/spotify/threaddump-analyzer>
 - It's JavaScript based and can be executed locally on a html page.
 - Threads with identical stacks are grouped
 - Locks between threads are visualized

Online Java Thread Dump Analyzer

Paste your thread dump here:

```
"Reference Handler" - Thread t@2
  java.lang.Thread.State: WAITING
    at java.lang.Object.wait(Native Method)
    - waiting on <5c019e1f> (a java.lang.ref.Reference$Lock)
    at java.lang.Object.wait(Object.java:502)
    at java.lang.ref.Reference$ReferenceHandler.run(Reference.java:157)

  Locked ownable synchronizers:
    - None
```

Analyze Clear

This page is client-side only, and no data will leave your computer when you click Analyze.

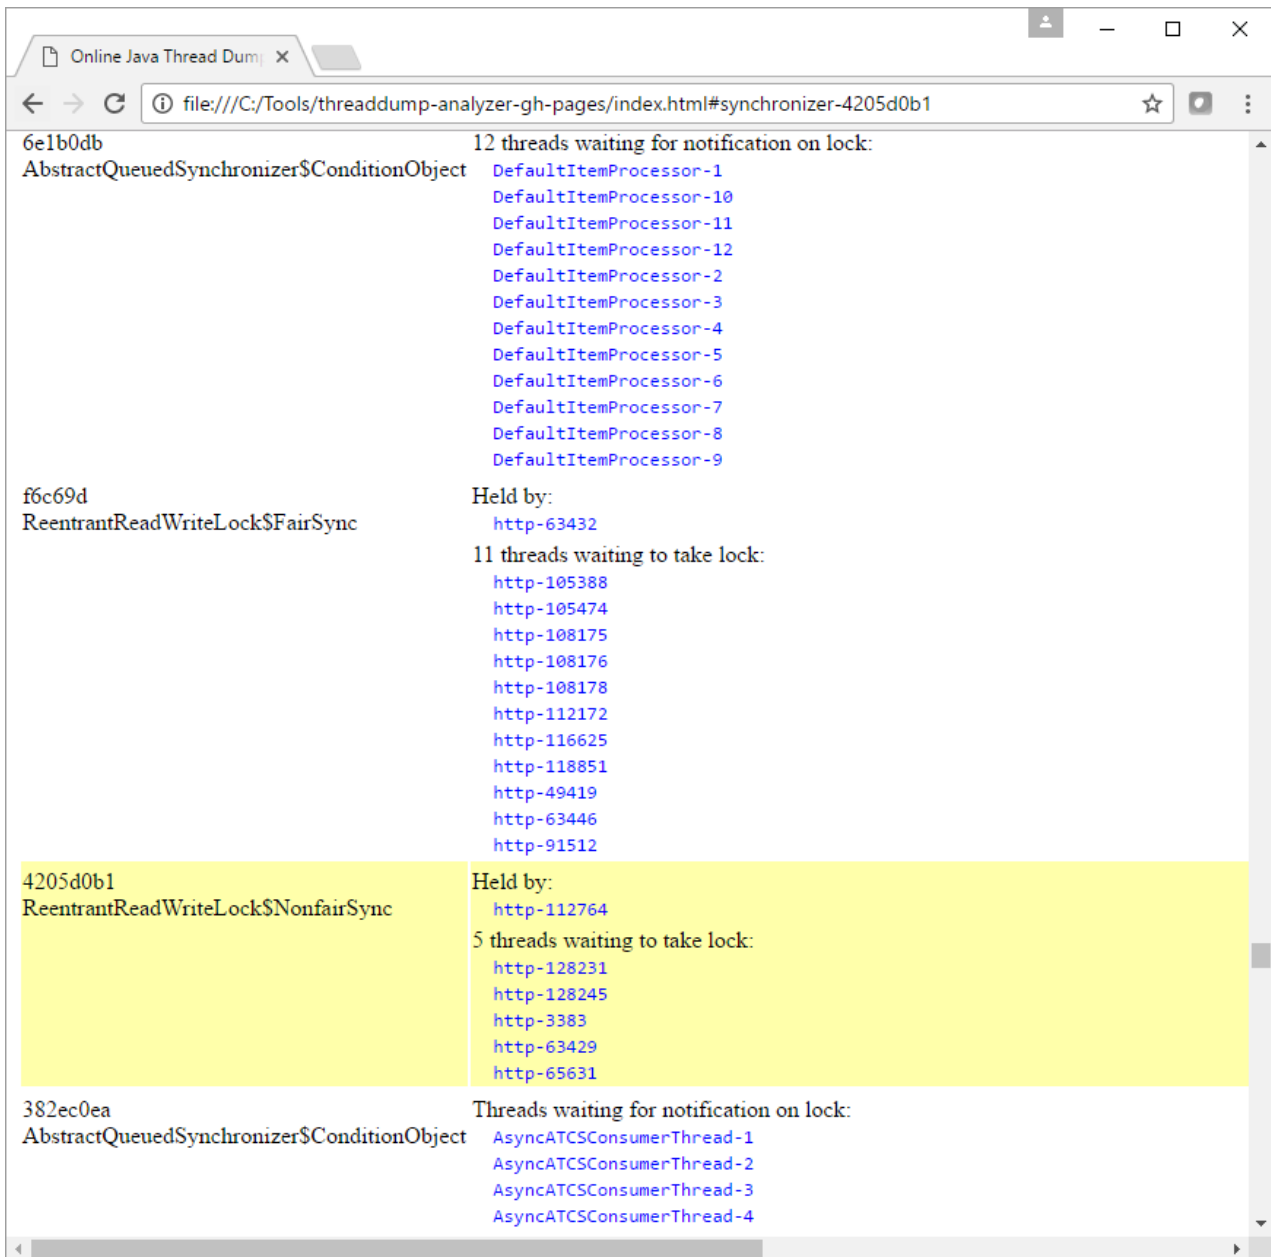
Top Methods From 71 Running Threads

```
55 java.net.SocketInputStream.socketRead0(Native Method)
 8 sun.nio.ch.WindowsSelectorImpl$SubSelector.poll0(Native Method)
 5 java.net.DualStackPlainSocketImpl.accept0(Native Method)
 1 java.io.FileInputStream.readBytes(Native Method)
 1 sun.management.ThreadImpl.dumpThreads0(Native Method)
 1 sun.nio.ch.ServerSocketChannelImpl.accept0(Native Method)
```

406 threads found

79 threads with this stack:

```
"AsyncATCSConsumerThread-1": awaiting notification on [382ec0ea]
"AsyncATCSConsumerThread-2": awaiting notification on [382ec0ea]
"AsyncATCSConsumerThread-3": awaiting notification on [382ec0ea]
"AsyncATCSConsumerThread-4": awaiting notification on [382ec0ea]
"CommunicationWorker-1": awaiting notification on [5799a2ec]
"CommunicationWorker-10": awaiting notification on [5799a2ec]
```

8 How to enable Java Management Extensions (JMX)

The Product 360 application server has various performance and management extensions which are available through the standard Java Management Extensions (JMX) interface. By default, the application server is started with the needed JMX parameters in place. This page describes those parameters for older versions or just as a reference.

Development and Test environment

To enable PIM - Server jmx you need to set jmx java properties. In the development environment use eclipse launch config. In test environment set them in the `console_debug.cmd`:

How to enable JMX local access

```
-Dcom.sun.management.jmxremote
```

8.1 Production environment

In production use `wrapper.conf` file:

Tanuki Wrapper

Because of bugs in the tanuki wrapper the jmv processes sometimes are not visible when started as a windows service (see Tanuki FAQ). In this case you need to configure jmx remote access. If you are running a production system make sure that firewall protects the jmx port from unauthorized access. See oracle JMX configuration manual for more details

How to enable JMX remote access

```
#This is not safe from security point of view!
-Dcom.sun.management.jmxremote.port=55555
-Dcom.sun.management.jmxremote.authenticate=false
-Dcom.sun.management.jmxremote.ssl=false
```

Add this lines to the `wrapper.conf` file in the Block # VM OPTIONS. Continue the existing lines (ending with "wrapper.java.additional.25"):

```
...
wrapper.java.additional.26 = -Dcom.sun.management.jmxremote.port=55555
wrapper.java.additional.27 = -Dcom.sun.management.jmxremote.authenticate=false
wrapper.java.additional.28 = -Dcom.sun.management.jmxremote.ssl=false
```

8.2 Download and Installation of JMX Tooling


To monitor managed beans you can use `jconsole.exe` (available in the JDK 1.5 and higher) or `visualvm` (available in the JDK 1.6 and higher or as a standalone application) with JMX plugins. Please note that the development department will need these tools on the corresponding machine if you want us to have a look on the system - it always streamlines the help processes if those tools are already installed.

User permissions

In some cases you will need to start `jconsole`/`VisualVM` with admin permissions (in windows file explorer "Start as Administrator") in order to see java process started as services

8.2.1 Connect on productive system

1. Open Visual VM
2. Goto File / Add JMX Connection
3. Enter localhost:55555 for the host, optionally add some display name for the connection



4. Press Ok
5. Connect to the application HPM Server by double click on the application name in the Applications view

Copyright

Copyright (c) 1993-2020 Informatica LLC. All rights reserved.

This Software and documentation contain proprietary information of Informatica LLC and are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright law. Reverse engineering of the software is prohibited. No part of this document may be reproduced or transmitted in any form, by any means (electronic, photocopying, recording or otherwise) without prior consent of Informatica LLC. This Software may be protected by [U.S. and/or](#) international Patents and other Patents Pending.

Use, duplication, or disclosure of the Software by the U.S. Government is subject to the restrictions set forth in the applicable software license agreement and as provided in DFARS 227.7202-1(a) and 227.7702-3(a) (1995), DFARS 252.227-7013©(1)(ii) (OCT 1988), FAR 12.212(a) (1995), FAR 52.227-19, or FAR 52.227-14

(ALT III), as applicable.

The information in this product or documentation is subject to change without notice. If you find any problems in this product or documentation, please report them to us in writing.

Informatica, Informatica Platform, Informatica Data Services, PowerCenter, PowerCenterRT, PowerCenter Connect, PowerCenter Data Analyzer, PowerExchange, PowerMart, Metadata Manager, Informatica Data Quality, Informatica Data Explorer, Informatica B2B Data Transformation, Informatica B2B Data Exchange Informatica On Demand, Informatica Identity Resolution, Informatica Application Information Lifecycle Management, Informatica Complex Event Processing, Ultra Messaging and Informatica Master Data Management are trademarks or registered trademarks of Informatica LLC in the United States and in jurisdictions throughout the world. All other company and product names may be trade names or trademarks of their respective owners.

Portions of this software and/or documentation are subject to copyright held by third parties, including without limitation: Copyright DataDirect Technologies. All rights reserved. Copyright © Sun Microsystems. All rights reserved. Copyright © RSA Security Inc. All Rights Reserved. Copyright © Ordinal Technology Corp. All rights reserved. Copyright © Aandacht c.v. All rights reserved. Copyright Genivia, Inc. All rights reserved. Copyright Isomorphic Software. All rights reserved. Copyright © Meta Integration Technology, Inc. All rights reserved. Copyright © Intalio. All rights reserved. Copyright © Oracle. All rights reserved. Copyright © Adobe Systems Incorporated. All rights reserved. Copyright © DataArt, Inc. All rights reserved. Copyright © ComponentSource. All rights reserved. Copyright © Microsoft Corporation. All rights reserved. Copyright © Rogue Wave Software, Inc. All rights reserved. Copyright © Teradata Corporation. All rights reserved. Copyright © Yahoo! Inc. All rights reserved. Copyright © Glyph & Cog, LLC. All rights reserved. Copyright © Thinkmap, Inc. All rights reserved. Copyright © Clearpace Software Limited. All rights reserved. Copyright © Information Builders, Inc. All rights reserved. Copyright © OSS Nokalva, Inc. All rights reserved. Copyright Edifecs, Inc. All rights reserved.

Copyright Cleo Communications, Inc. All rights reserved. Copyright © International Organization for Standardization 1986. All rights reserved. Copyright © ej-technologies GmbH. All rights reserved. Copyright © Jaspersoft Corporation. All rights reserved. Copyright © International Business Machines Corporation. All rights reserved. Copyright © yWorks GmbH. All rights reserved. Copyright © Lucent Technologies. All rights reserved. Copyright (c) University of Toronto. All rights reserved. Copyright © Daniel Veillard. All rights reserved. Copyright © Unicode, Inc. Copyright IBM Corp. All rights reserved. Copyright © MicroQuill Software Publishing, Inc. All rights reserved. Copyright © PassMark Software Pty Ltd. All rights reserved. Copyright © LogiXML, Inc. All rights reserved. Copyright © 2003-2010 Lorenzi Davide, All rights reserved. Copyright © Red Hat, Inc. All rights reserved. Copyright © The Board of Trustees of the Leland Stanford Junior University. All rights reserved. Copyright

© EMC Corporation. All rights reserved. Copyright © Flexera Software. All rights reserved. Copyright © Jinfonet Software. All rights reserved. Copyright © Apple Inc. All rights reserved. Copyright © Telerik Inc. All rights reserved. Copyright © BEA Systems. All rights reserved. Copyright © PDFlib GmbH. All rights reserved. Copyright © Orientation in Objects GmbH. All rights reserved. Copyright © Tanuki Software, Ltd. All rights reserved. Copyright © Ricebridge. All rights reserved. Copyright © Sencha, Inc. All rights reserved. Copyright © Scalable Systems, Inc. All rights reserved. Copyright © jqWidgets. All rights reserved. Copyright © Tableau Software, Inc. All rights reserved. Copyright © MaxMind, Inc. All Rights Reserved. Copyright © TMate Software s.r.o. All rights reserved. Copyright © MapR Technologies Inc. All rights reserved.

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>), and/or other software which is licensed under various versions of the Apache License (the "License"). You may obtain a copy of these Licenses at <http://www.apache.org/licenses/>. Unless required by applicable law or agreed to in writing, software distributed under these Licenses is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the Licenses for the specific language governing permissions and limitations under the Licenses.

This product includes software which was developed by Mozilla (<http://www.mozilla.org/>), software copyright The JBoss Group, LLC, all rights reserved; software copyright © 1999-2006 by Bruno Lowagie and Paulo Soares and other software which is licensed under various versions of the GNU Lesser General Public License Agreement, which may be found at [http:// www.gnu.org/licenses/lgpl.html](http://www.gnu.org/licenses/lgpl.html). The materials are provided free of charge by Informatica, "as-is", without warranty of any kind, either express or implied, including but not limited to the implied warranties of merchantability and fitness for a particular purpose.

The product includes ACE(TM) and TAO(TM) software copyrighted by Douglas C. Schmidt and his research group at Washington University, University of California, Irvine, and Vanderbilt University, Copyright (©) 1993-2006, all rights reserved.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (copyright The OpenSSL Project. All Rights Reserved) and redistribution of this software is subject to terms available at <http://www.openssl.org> and <http://www.openssl.org/source/license.html>.

This product includes Curl software which is Copyright 1996-2013, Daniel Stenberg, < daniel@haxx.se >. All Rights Reserved. Permissions and limitations regarding this software are subject to terms available at [http://curl.haxx.se/docs/ copyright.html](http://curl.haxx.se/docs/copyright.html). Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

The product includes software copyright 2001-2005 (©) MetaStuff, Ltd. All Rights Reserved. Permissions and limitations regarding this software are subject to terms available at [http://www.dom4j.org/ license.html](http://www.dom4j.org/license.html).

The product includes software copyright © 2004-2007, The Dojo Foundation. All Rights Reserved. Permissions and limitations regarding this software are subject to terms available at [http://dojotoolkit.org/ license](http://dojotoolkit.org/license).

This product includes ICU software which is copyright International Business Machines Corporation and others. All rights reserved. Permissions and limitations regarding this software are subject to terms available at <http://source.icu-project.org/repos/icu/icu/trunk/license.html>.

This product includes software copyright © 1996-2006 Per Bothner. All rights reserved. Your right to use such materials is set forth in the license which may be found at [http://www.gnu.org/software/kawa/ Software-License.html](http://www.gnu.org/software/kawa/Software-License.html).

This product includes OSSP UUID software which is Copyright © 2002 Ralf S. Engelschall, Copyright © 2002 The OSSP Project Copyright © 2002 Cable & Wireless Deutschland. Permissions and limitations regarding this software are subject to terms available at <http://www.opensource.org/licenses/mit-license.php>.

This product includes software developed by Boost (<http://www.boost.org/>) or under the Boost software license. Permissions and limitations regarding this software are subject to terms available at [http:// www.boost.org/LICENSE_1_0.txt](http://www.boost.org/LICENSE_1_0.txt).

This product includes software copyright © 1997-2007 University of Cambridge. Permissions and limitations regarding this software are subject to terms available at <http://www.pcre.org/license.txt>. This product includes software copyright © 2007 The Eclipse Foundation. All Rights Reserved. Permissions and limitations regarding this software are subject to terms available at [http:// www.eclipse.org/org/documents/epl-v10.php](http://www.eclipse.org/org/documents/epl-v10.php) and at <http://www.eclipse.org/org/documents/edl-v10.php>.

This product includes software licensed under the terms at <http://www.tcl.tk/software/tcltk/license.html>, <http://www.bosrup.com/web/overlib/?License>, [http:// www.stlport.org/doc/ license.html](http://www.stlport.org/doc/license.html), <http://asm.ow2.org/license.html>, <http://www.cryptix.org/LICENSE.TXT>, <http://hsqldb.org/web/hsqLicense.html>, <http://unit.sourceforge.net/doc/license.html>, <http://jung.sourceforge.net/license.txt> , http://www.gzip.org/zlib/zlib_license.html, [http://www.openldap.org/software/release/ license.html](http://www.openldap.org/software/release/license.html), <http://www.libssh2.org>, <http://slf4j.org/license.html>, <http://www.sente.ch/software/OpenSourceLicense.html>, [http://fusesource.com/downloads/license- agreements/fuse-message-broker-v-5-3- licenseagreement](http://fusesource.com/downloads/license-agreements/fuse-message-broker-v-5-3-licenseagreement); [http:// antlr.org/license.html](http://antlr.org/license.html); <http://aopalliance.sourceforge.net/>; [http://www.bouncycastle.org/ licence.html](http://www.bouncycastle.org/licence.html); <http://www.jgraph.com/jgraphdownload.html>; <http://www.jcraft.com/jsch/LICENSE.txt>; http://jotm.objectweb.org/bsd_license.html; . [http://www.w3.org/Consortium/Legal/ 2002/copyright-](http://www.w3.org/Consortium/Legal/2002/copyright-)

software-20021231; <http://www.slf4j.org/license.html>; <http://nanoxml.sourceforge.net/orig/copyright.html>; <http://www.json.org/license.html>; <http://forge.ow2.org/projects/javaservice/>, <http://www.postgresql.org/about/licence.html>, <http://www.sqlite.org/copyright.html>, <http://www.tcl.tk/software/tcltk/license.html>, <http://www.jaxen.org/faq.html>, <http://www.jdom.org/docs/faq.html>, <http://www.slf4j.org/license.html>; <http://www.iodbc.org/dataspace/iodbc/wiki/iODBC/License>; <http://www.keplerproject.org/md5/license.html>; <http://www.toedter.com/en/jcalendar/license.html>; <http://www.edankert.com/bounce/index.html>; <http://www.net-snmp.org/about/license.html>; <http://www.openmdx.org/#FAQ>; http://www.php.net/license/3_01.txt; <http://srp.stanford.edu/license.txt>; <http://www.schneier.com/blowfish.html>; <http://www.jmock.org/license.html>; <http://xsom.java.net>; <http://benalman.com/about/license/>; <https://github.com/CreateJS/EaselJS/blob/master/src/easeljs/display/Bitmap.js>; <http://www.h2database.com/html/license.html#summary>; <http://jsoncpp.sourceforge.net/LICENSE>; <http://jdbc.postgresql.org/license.html>; <http://protobuf.googlecode.com/svn/trunk/src/google/protobuf/descriptor.proto>; <https://github.com/rantav/hector/blob/master/LICENSE>; <http://web.mit.edu/Kerberos/krb5-current/doc/mitK5license.html>; <http://jibx.sourceforge.net/jibx-license.html>; <https://github.com/lyokato/libgeohash/blob/master/LICENSE>; <https://github.com/hjiang/jsonxx/blob/master/LICENSE>; <https://code.google.com/p/lz4/>; <https://github.com/jedisct1/libsodium/blob/master/LICENSE>; <http://one-jar.sourceforge.net/index.php?page=documents&file=license>; <https://github.com/EsotericSoftware/kryo/blob/master/license.txt>; <http://www.scala-lang.org/license.html>; <https://github.com/tinkerpop/blueprints/blob/master/LICENSE.txt>; and <http://gee.cs.oswego.edu/dl/classes/EDU/oswego/cs/dl/util/concurrent/intro.html>.

This product includes software licensed under the Academic Free License (<http://www.opensource.org/licenses/afl-3.0.php>), the Common Development and Distribution License (<http://www.opensource.org/licenses/cddl1.php>) the Common Public License (<http://www.opensource.org/licenses/cpl1.0.php>), the Sun Binary Code License Agreement Supplemental License Terms, the BSD License (<http://www.opensource.org/licenses/bsd-license.php>), the new BSDLicense (<http://opensource.org/licenses/BSD-3-Clause>), the MIT License (<http://www.opensource.org/licenses/mitlicense.php>), the Artistic License (<http://www.opensource.org/licenses/artistic-license-1.0>) and the Initial Developer's Public License Version 1.0 (<http://www.firebirdsql.org/en/initial-developer-s-public-license-version-1-0/>).

This product includes software copyright © 2003-2006 Joe Walnes, 2006-2007 XStream Committers. All rights reserved. Permissions and limitations regarding this software are subject to terms available at <http://xstream.codehaus.org/license.html>. This product includes software developed by the Indiana University Extreme! Lab. For further information please visit <http://www.extreme.indiana.edu/>.

This product includes software Copyright (c) 2013 Frank Balluffi and Markus Moeller. All rights reserved. Permissions and limitations regarding this software are subject to terms of the MIT license.

This Software is protected by U.S. Patent Numbers 5,794,246; 6,014,670; 6,016,501; 6,029,178; 6,032,158; 6,035,307; 6,044,374; 6,092,086; 6,208,990; 6,339,775; 6,640,226; 6,789,096; 6,823,373; 6,850,947; 6,895,471; 7,117,215; 7,162,643; 7,243,110; 7,254,590; 7,281,001; 7,421,458; 7,496,588; 7,523,121; 7,584,422; 7,676,516; 7,720,842; 7,721,270; 7,774,791; 8,065,266; 8,150,803; 8,166,048; 8,166,071; 8,200,622; 8,224,873; 8,271,477; 8,327,419; 8,386,435; 8,392,460; 8,453,159; 8,458,230; 8,707,336; 8,886,617 and RE44,478, International Patents and other Patents Pending.

DISCLAIMER: Informatica LLC provides this documentation "as is" without warranty of any kind, either express or implied, including, but not limited to, the implied warranties of noninfringement, merchantability, or use for a particular purpose. Informatica LLC does not warrant that this software or documentation is error free. The information provided in this software or documentation may include technical inaccuracies or typographical errors. The information in this software and documentation is subject to change at any time without notice.iv>

NOTICES

This Informatica product (the "Software") includes certain drivers (the "DataDirect Drivers") from DataDirect Technologies, an operating company of Progress Software Corporation ("DataDirect") which are subject to the following terms and conditions:

1. THE DATADIRECT DRIVERS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT.
2. IN NO EVENT WILL DATADIRECT OR ITS THIRD PARTY SUPPLIERS BE LIABLE TO THE END-USER CUSTOMER FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR OTHER DAMAGES ARISING OUT OF THE USE OF THE ODBC DRIVERS, WHETHER OR NOT INFORMED OF THE POSSIBILITIES OF DAMAGES IN ADVANCE. THESE LIMITATIONS APPLY TO ALL CAUSES OF ACTION, INCLUDING, WITHOUT LIMITATION, BREACH OF CONTRACT, BREACH OF WARRANTY, NEGLIGENCE, STRICT LIABILITY,

MISREPRESENTATION AND OTHER TORTS.