



*cutting through complexity™*

ADVISORY

**CGI**

Sm@rtering Volume Tests

Monitoring report

**kpmg.pt**

June 2015

# Contents

|                                      |    |
|--------------------------------------|----|
| Executive Summary                    | 1  |
| Test Setup                           | 3  |
| Test Scenarios                       | 7  |
| Test Results                         | 9  |
| References and Support documentation | 11 |
| Volume Tests Physical Infrastructure | 12 |

KPMG's conclusions are made specifically under the conditions described in this report. Furthermore, the projection of any conclusions for future periods is subject to the risk that changes made to the system, or the failure to make needed changes to the system, may alter the validity of such conclusions. No one should act on such information without appropriate professional advice after a thorough examination of the particular situation.

The information contained herein is solely for the benefit of the parties identified in our engagement letter and should not be copied, cited or referred, in whole or in part, without our previous written consent. We decline any responsibility regarding any other party to whom this information may be provided, or that may acquire a copy.

© 2015 KPMG Advisory - Consultores de Gestão, S.A., a Portuguese company and a member firm of the KPMG network of independent member firms affiliated with KPMG International Cooperative ("KPMG International"), a Swiss entity. All rights reserved.

## Executive Summary

The Sm@rtering platform allows for the collection, management and processing of data from Advanced Metering Infrastructures (AMI). CGI performed Volume Testing procedures, monitored by KPMG, to validate the ability of the Sm@rtering platform to process millions of meters. KPMG has confirmed that the Sm@rtering platform successfully processed all data relative to 5 million meters, for all the business processes within the test scope.

The Sm@rtering platform is a modular system, based on three modules:

- **AMI HEAD-END**, which includes all the functionalities that allow scheduling and communicating with the AMI network (e.g. meters and sensors), using different types of protocols. This component is also responsible for the translation of the data to a standard Sm@rtering format and for loading data into the system.
- **MDM/EDM**, holds all the business logic necessary to support the installation processes and operation of the metering equipment, including the storage, validation and the estimation of all the data collected. This module also includes the ability to calculate billing determinants that provide information to other business processes (e.g. billing).
- **AMI SUPERVISION & GRID (SUP)**, allows the monitoring and control of the performance of the AMI and also monitoring of the status of the electrical grid and correlation with other relevant data to anticipate and reduce the outages. This module is supported by operational dashboards and real-time event processing.

The main goal of these Volume Tests is to demonstrate the ability of the different Sm@rtering modules to manage the data of a relevant number of meters. The Volume Testing was performed using release 3.5 of the Sm@rtering Platform.

Considering its criticality, CGI designed and performed 8 test scenarios to evaluate the capability of the Sm@rtering platform to handle the data generated by either 5 million meters or 10 million meters (depending on the test scenario).

KPMG monitored the CGI team during the execution of the test scenarios to (a) confirm that the test scenarios were executed in accordance with the predefined workload and test conditions; (b) ensure that the results were monitored, recorded and reported accurately; and (c) verify if the results met the success criteria.

In this exercise, KPMG procedures were focused on monitoring and reporting the data volumes used and the time required to execute each test scenario. The measure of success was the creation/update of collected/processed data in the relevant database tables. KPMG did not perform procedures to validate the Sm@rtering platform business logic, further evaluate the Sm@rtering platform performance, or assess the correctness of the processed data.

The results shown in the following table were recorded by KPMG:

| Test Scenarios |  |   |  |
|----------------|--|---|--|
| Scenario       | Name   | Business Processes executed   | Elapsed time until last write in the database (hh:mm:ss) |
| 1              | Daily readings, 5 million meters                       | 1. Load Assets Data<br>2. Validate, Edit and Estimate<br>3. Billing Determinants                        | 01:22:58   |
| 2              | 1h Load Profiles, 5 million meters                     | 1. Load Assets Data<br>2. Validate, Edit and Estimate   | 01:24:37   |
| 3              | Real Time Events, 5 million meters                     | 4. Event Processing   | 00:47:30   |
| 4              | Data Collection, All, 5 million meters                 | 1. Load Assets Data<br>2. Validate, Edit and Estimate<br>3. Billing Determinants<br>4. Event Processing | 04:36:30   |
| 5              | Calculate Key Performance Indicators, 5 million meters | 5. Calculate KPIs / Dashboards  | 00:34:24   |
| 6              | 15min Load Profiles, 5 million meters                  | 1. Load Assets Data<br>2. Validate, Edit and Estimate   | 01:56:12   |
| 7              | 1h Load Profiles, 10 million meters                    | 1. Load Assets Data<br>2. Validate, Edit and Estimate   | 02:55:20   |
| 8              | Real Time Events, 10 million meters                    | 4. Event Processing   | 01:41:49   |

The test results demonstrate that the Sm@rtering platform is capable of processing the data volumes described in each of the test scenarios in the reported elapsed times.

The test scenarios were performed under specific conditions. Those specific conditions and assumptions considered for each business process are detailed in the "Test Setup" chapter. KPMG conclusions are made specifically under these conditions. Furthermore, the projection of any conclusions to future periods is subject to the risk that changes made to the system, or the failure to make required changes to the system, may alter the validity of such conclusions.

## Test Setup

The Sm@rtering platform allows for the collection, management and processing data from Advanced Metering Infrastructures (AMI). CGI performed Volume Testing procedures, monitored by KPMG, to validate the ability of the Sm@rtering platform to process the data generated by either 5 or 10 million meters (depending on the test scenario).

The Sm@rtering platform is a modular system, based on three modules:

- **AMI HEAD-END**, which includes all the functionalities that allow scheduling and communicating with the AMI network (e.g. meters and sensors), using different types of protocols. This component is also responsible for the translation of the data to a standard Sm@rtering format and for loading data into the system.
- **MDM/EDM**, holds all the business logic necessary to support the installation processes and operation of the metering equipment, including the storage, validation and the estimation of all the data collected. This module also includes the ability to calculate billing determinants that provide information to other business processes (e.g. billing).
- **AMI SUPERVISION & GRID (SUP)**, allows the monitoring and control of the performance of the AMI and also monitoring of the status of the electrical grid and correlation with other relevant data to anticipate and reduce the outages. This module is supported by operational dashboards and real-time event processing.

The main goal of these Volume Tests is to demonstrate the capability of the different Sm@rtering modules to manage the data of a relevant number of meters. The Volume Testing was performed using release 3.5 of the Sm@rtering platform.

## METHODOLOGY

The tests were oriented to **business processes** that represent normal processing in a working day.



The table below describes the methodology followed to test each of these **business processes** and specific configurations made for volume testing purposes:

| Business Processes                                |  |
|---|--|
| Name  | Description  |
| <b>Load Assets Data</b><br>(Head-End)             | <p>This process is responsible for the processing, translation of the data to the standard Sm@rtering format and loading onto the Database.</p> <p>The Load Assets Data process is event driven and therefore executed automatically as soon as data is available for processing (i.e. the file with the data is available for processing).</p> <p>There are three types of data being generated, depending on the test scenario:</p> <ul style="list-style-type: none"> <li>▪ Daily register read files which include 4 records per registered meter</li> <li>▪ Load profile files of active energy in periods of 1h (24 records per day)</li> <li>▪ Load profile files of active energy in periods of 15min (96 records per day)</li> </ul> <p>This process ends when all the data is updated in the database.</p> <p>These test scenarios don't consider the "communication" time between the AMI network and the Sm@rtering system. They assume that all data is readily available.</p>  |
| <b>VEE – Validate, Edit and Estimate</b><br>(MDM) | <p>This process validates the existence of meter data (daily register reads and load profiles) and checks if it is valid or not. If needed, it calculates consumption estimates based on a set of predefined rules for the type of client and the type of data being processed.</p> <p>The Validation process is event driven and therefore executed automatically as the data is available on the database.</p> <p>Below is the list of validations performed by this process:</p> <ul style="list-style-type: none"> <li>▪ Daily Readings               <ul style="list-style-type: none"> <li>▪ Meter read is zero</li> <li>▪ Number of registers is different from expected</li> <li>▪ Number of digits</li> <li>▪ Negative delta read</li> <li>▪ Meter read is outside of expected interval</li> </ul> </li> <li>▪ Load Profile               <ul style="list-style-type: none"> <li>▪ Missing data</li> <li>▪ Invalid data status</li> <li>▪ Load profiles <i>versus</i> daily reads</li> <li>▪ 25 Hours</li> </ul> </li> </ul> <p>All invalid data is registered with a specific status. The validations that failed are later analyzed by the estimation mechanism which tries to populate the missing/invalid data according to the rules set, based on the meter history.</p> <p>The process ends when all the data is loaded or estimated.</p> <p>The rate of invalid data used in these tests was reported by CGI to be 2%.</p> <p>Note: The impact on performance due to complexity of the estimations algorithm was not evaluated.</p> |

## Business Processes

| Name  | Description  |
|---|--|
| <b>Billing Determinants</b><br>(MDM)        | <p>The calculation processes are important to support various business processes, including billing.</p> <p>This process is event driven and is executed automatically as the VEE – Validate, Edit and Estimate process is finished.</p> <p>This process ends when the consumption for all meters (with data) is calculated.</p> <p>For the purpose of these tests only the calculation process that calculates consumptions from the daily register read was used.</p>  |
| <b>Events Processing</b><br>(SUP)           | <p>This process is responsible for receiving, processing and registering all the events from the AMI infrastructure and other support systems. The process handles the events and has filtering, enrichment and correlation mechanisms. It also creates other events depending on the data received and the configured rules.</p> <p>This process is event driven and is executed automatically as soon as events are received (i.e. the file with the records of events is available for processing).</p> <p>During the Volume Tests, several types of events will be considered, including:</p> <ul style="list-style-type: none"> <li>▪ Power Outage</li> <li>▪ Network quality</li> </ul> <p>For these tests, a single file with all the events was used, which stresses the test to the fullest, as in real life not all daily events happen at the same time. Those events were processed in parallel.</p> |
| <b>Calculate KPIs / Dashboards</b><br>(SUP) | <p>This process uses data from the previous processes to calculate the performance indicators that will feed the various monitoring and control processes (Dashboards, AMI Status, Benchmark).</p> <p>This process calculates the KPIs for all meters and then aggregates those values into different levels of hierarchy, thus supporting the top/down view that Sm@rtering provides in its Dashboards.</p> <p>This process is prepared to calculate in “batch mode” indicators that vary little through time, and in “real time” for indicators that vary more frequently.</p> <p>For Volume Test purposes, only the KPIs related to the test processes (e.g. loading, validation, event processing) will be calculated and in “batch mode”.</p>   |

As mentioned, the business processes are event driven, that is, they start automatically as soon as they have inputs for processing. For the purpose of these tests, CGI developed **support processes** in order to automate the test tasks. These support processes include:

- Database management procedures (e.g. initial database loading, restoring initial conditions, clean test support tables)
- AMI data simulation (e.g. daily register read, meter interval data, events)
- Automation of the different Test Scenarios
- Data collection at the end of each Scenario

The initial database setup included all the meters and master data required for the test scenarios (5M or 10M meters). All the data loaded in the database during the test scenarios has been kept in the database. At the end of the execution of each scenario, a data collection task retrieves all relevant data and produces a final report with the test counters and execution times.

KPMG evaluated the support processes including how they work, how the number of processed records is assessed and how time is measured.

CGI has adopted the following steps in order to implement the Volume Tests:

1. **INSTALL SM@RTERING PLATFORM** in the test infrastructure;
2. **PREPARE PROCESSES AND DATA.** Create the data volumes required on each scenario and simulate the data sent by the AMI infrastructure (support processes) in order to start the Sm@rtering event driven processes;
3. **PREPARE TEST SCENARIOS**, which includes running the necessary support procedures to establish the initial test conditions for each scenario;
4. **EXECUTE TESTS**, and record test results;
5. **REPORT TEST RESULTS.**

After the execution of each test scenario, there is a support process that queries the DB and calculates the number of records processed and the execution time. This report provides the volume of data that was processed during the test and the time in which it was executed.

The test execution scripts retrieve the initial and ending timestamp for each test, taking into consideration the actions being performed. As a rule of thumb, the initial timestamp is recorded as soon as the data is available to the system. The ending timestamp is recorded once the last write is done in the database.

There is also a support process that checks for the termination conditions (e.g. the existence of new data in the Queues, active threads, data written in the database) which runs every 10 minutes. Only after making sure that all processes and threads are inactive and that there is no new data to process, the system can detect that everything is completed. When this happens, there is one additional entry in the log and the execution script terminates. KPMG decided to also report the times until the timestamp of this log entry (called "sanity check"), which is the last moment at which any processing related to an individual test could have been done.

In the Test Results section of this document, KPMG presents both values, the last write in the database and the sanity check times.

## Test Scenarios

The Sm@rtering Volume Test scenarios include the most relevant business processes that occur on a normal working day, based on expected distributions of the data.

The Volume Test methodology was focused on replicating the normal processes and data that are expected in one business day, based on different data processing strategies (e.g. parallel processing, sequential processing).

The Volume Testing focused on two main objectives:

- Validation of the base processes for all modules in the Sm@rtering Platform for a volume of 5 Million meters, simulating all the normal processes in a normal working day;
- Validation of the critical collection processes (events and meter interval data) for a volume of 10 Million meters.

The two main objectives were then divided into eight scenarios that simulates the behaviour of the Sm@rtering platform in different operating scenarios.

The table below describes the eight scenarios that were designed:

| Test Scenarios |                          |   |  |   |   |  |
|----------------|--------------------------|---|--|---|---|--|
| Scenario       | Name                     | Volume  | Type of Data   | Business Processes  | Description   | Process execution details                                    |
| 1              | Daily readings, 5M       | 5M Meters<br>20M Daily register reads                                   | Daily register reads                                       | 1. Load Assets Data<br>2. Validate, Edit and Estimate<br>3. Billing Determinants                        | Collection of daily register reads, processing and validation. Calculation of billing determinants.   | Processes were executed sequentially by meter (event driven) |
| 2              | 1h Load Profiles, 5M     | 5M Meters<br>120M Time Series   | Meter interval data (1h)                                   | 1. Load Assets Data<br>2. Validate, Edit and Estimate   | Collection of meter interval data, processing and validation  | Processes were executed sequentially by meter                |
| 3              | Real Time Events, 5M     | 5M Meters<br>15M Events   | Events   | 4. Event Processing   | Collection, processing and validation of events in real time  |  |
| 4              | Data Collection, All, 5M | 5M Meters<br>20M Daily register reads<br>120M Time Series<br>15M Events | Daily register reads<br>Meter interval data (1h)<br>Events | 1. Load Assets Data<br>2. Validate, Edit and Estimate<br>3. Billing Determinants<br>4. Event Processing | Simultaneous collection, processing, validation and calculation of billing determinants for daily register reads, meter interval data and events. | Processes were executed in parallel                          |
| 5              | Calculate KPIs, 5M       | 5M Meters   | Own database records                                       | 5. Calculate KPIs / Dashboards  | KPI processing  | Processes were executed in "batch mode"                      |

## Test Scenarios

| Scenario | Name                    | Volume                         | Type of Data                | Business Processes                                    | Description   | Process execution details                     |
|----------|-------------------------|--------------------------------|-----------------------------|---|---|---|
| 6        | 15min Load Profiles, 5M | 5M Meters<br>480M Time Series  | Meter interval data (15min) | 1. Load Assets Data<br>2. Validate, Edit and Estimate | Simultaneous collection, processing, validation and calculation of billing determinants for daily register reads, meter interval data and events. | Processes were executed sequentially by meter |
| 7        | 1h Load Profiles, 10M   | 10M Meters<br>240M Time Series | Meter interval data (1h)    | 1. Load Assets Data<br>2. Validate, Edit and Estimate | Collection of meter interval data, processing and validation  | Processes were executed sequentially by meter |
| 8        | Real Time Events, 10M   | 10M Meters<br>30M Events       | Events                      | 4. Event Processing                                   | Collection, processing and validation of events in real time  |   |

## Test Results

The Sm@rtering Volume Tests were executed under the supervision of KPMG. Results were recorded for each of the scenarios depicted in the previous section, taking into consideration the volume of data and the time Sm@rtering took to process those volumes of data.

KPMG accompanied CGI in the execution of each of the test scenarios, according to CGI's definition as described in the previous section.

The following results were recorded for each test Scenario:

| Test Scenarios |                      |                                       |  |   |
|----------------|----------------------|---------------------------------------|--|---|
| Scenario Name  | Volume               | Business Processes                    | Description  | Results (hh:mm:ss)  |
| 1              | Daily readings, 5M   | 5M Meters<br>20M Daily register reads | 1. Load Assets Data<br>2. Validate, Edit and Estimate<br>3. Billing Determinants | <p>Collection of daily register reads, processing and validation. Calculation of billing determinants.</p> <p><b>Last write method</b><br/>01:22:58<br/>4.018 rec/sec</p> <p><b>Sanity check</b><br/>01:24:27<br/>3.947 rec/sec</p> |
| 2              | 1h Load Profiles, 5M | 5M Meters<br>120M Time Series         | 1. Load Assets Data<br>2. Validate, Edit and Estimate                            | <p>Collection of meter interval data, processing and validation</p> <p><b>Last write method</b><br/>01:24:37<br/>23.636 rec/sec</p> <p><b>Sanity check</b><br/>01:31:57<br/>21.751</p>  |
| 3              | Real Time Events, 5M | 5M Meters<br>15M Events               | 4. Event Processing  | <p>Collection, processing and validation of events in real time</p> <p><b>Last write method</b><br/>00:47:30<br/>5.263 rec/sec</p> <p><b>Sanity check</b><br/>00:48:01<br/>5.207 rec/sec</p>  |

| Test Scenarios |                          |   |   |  |
|----------------|--------------------------|---|---|--|
| Scenario Name  | Volume                   | Business Processes  | Description   | Results (hh:mm:ss)   |
| 4              | Data Collection, All, 5M | 5M Meters<br>20M Daily register reads<br>120M Time Series<br>15M Events | 1. Load Assets Data<br>2. Validate, Edit and Estimate<br>3. Billing Determinants<br>4. Event Processing | Simultaneous collection, processing, validation and calculation of billing determinants for daily register reads, meter interval data and events.<br><br><b>Last write method</b><br>04:36:30<br>9.343 rec/sec<br><br><b>Sanity check</b><br>04:40:07<br>9.222 rec/sec   |
| 5              | Calculate KPIs, 5M       | 5M Meters   | 5. Calculate KPIs / Dashboards  | KPI processing<br><br><b>Last write method</b><br>00:34:24<br>33.195 rec/sec<br><br><b>Sanity check</b><br>01:04:18<br>18.144 rec/sec  |
| 6              | 15min Load Profiles, 5M  | 5M Meters<br>480M Time Series   | 1. Load Assets Data<br>2. Validate, Edit and Estimate   | Simultaneous collection, processing, validation and calculation of billing determinants for daily register reads, meter interval data and events.<br><br><b>Last write method</b><br>01:56:12<br>68.847 rec/sec<br><br><b>Sanity check</b><br>02:03:29<br>64.786 rec/sec |
| 7              | 1h Load Profiles, 10M    | 10M Meters<br>240M Time Series  | 1. Load Assets Data<br>2. Validate, Edit and Estimate   | Collection of meter interval data, processing and validation<br><br><b>Last write method</b><br>02:55:20<br>22.814 rec/sec<br><br><b>Sanity check</b><br>03:16:54<br>20.315 rec/sec  |
| 8              | Real Time Events, 10M    | 10M Meters<br>30M Events  | 4. Event Processing   | Collection, processing and validation of events in real time<br><br><b>Last write method</b><br>01:41:49<br>4.911 rec/sec<br><br><b>Sanity check</b><br>01:50:57<br>4.507 rec/sec  |

The test results demonstrate that the Sm@rtering platform can process the volume of data described in each of the scenarios, within the reported time frames.

## References and supporting documentation

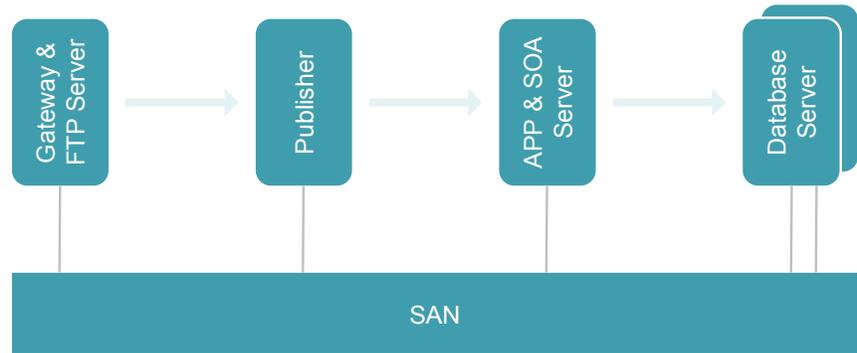
CGI provided several documents and other information that was used during the Volume Tests. The information provided included:

- R3.5 Sm@rtering platform - Volume Tests – Functional Specification v2.0
- Sm@rtering Volume Tests Kick-off documentation
- Various emails, including scenario scripts, table descriptions, doubts and clarifications.

The information described in these documents was the basis for the assumptions and calculations made by KPMG. CGI has agreed that all results are valid as per the information in these documents.

## Volume Tests Physical Infrastructure

CGI has implemented a dedicated infrastructure to support the Volume Tests:



The infrastructure includes the following servers:

| Volume Tests Infrastructure |   |
|-----------------------------|---|
| Name                        | Specifications  |
| Gateway & FTP Server        | Intel(R) Xeon(R) CPU E5-2609 v2 @ 2.50GHz<br>8 CPUs<br>4 cores/CPU<br>10 MB Cache<br>47 Gb Memory<br>136 GB HDD (local)<br>Linux version 2.6.32-431.el6.x86_64 (Red Hat 4.4.7-4)  |
| Publisher                   | Intel(R) Xeon(R) CPU E5-2650 v2 @ 2.60GHz<br>32 CPUs<br>8 cores/CPU<br>20 MB Cache<br>94 Gb Memory<br>136 GB HDD (local)<br>Linux version 2.6.32-431.el6.x86_64 (Red Hat 4.4.7-4) |
| Application & SOA Server    | Intel(R) Xeon(R) CPU E5-2650 v2 @ 2.60GHz<br>32 CPUs<br>8 cores/CPU<br>20 MB Cache<br>94 Gb Memory<br>136 GB HDD (local)<br>Linux version 2.6.32-431.el6.x86_64 (Red Hat 4.4.7-4) |

## Volume Tests Infrastructure

| Name                 | Specifications  |
|----------------------|---|
| Database Server (x2) | Intel(R) Xeon(R) CPU E5-2650 v2 @ 2.60GHz<br>32 CPUs<br>8 cores/CPU<br>20 MB Cache<br>94 Gb Memory<br>136 GB HDD (local)<br>Linux version 2.6.32-431.el6.x86_64 (Red Hat 4.4.7-4) |

## Contact us

**Gustavo Mendes**  
**Associate Partner**  
**IT Advisory**

**T** +351 210 110 012

**E** [gmendes@kpmg.com](mailto:gmendes@kpmg.com)

**Duarte Nóbrega**  
**Senior Manager**  
**IT Advisory**

**T** +351 210 110 012

**E** [dnobrega@kpmg.com](mailto:dnobrega@kpmg.com)

[www.kpmg.pt](http://www.kpmg.pt)

KPMG's conclusions are made specifically under the conditions described in this report. Furthermore, the projection of any conclusions for future periods is subject to the risk that changes made to the system, or the failure to make needed changes to the system, may alter the validity of such conclusions. No one should act on such information without appropriate professional advice after a thorough examination of the particular situation.

The information contained herein is solely for the benefit of the parties identified in our engagement letter and should not be copied, cited or referred, in whole or in part, without our previous written consent. We decline any responsibility regarding any other party to whom this information may be provided, or that may acquire a copy.

© 2015 KPMG Advisory - Consultores de Gestão, S.A., a Portuguese company and a member firm of the KPMG network of independent member firms affiliated with KPMG International Cooperative ("KPMG International"), a Swiss entity. All rights reserved. The KPMG name, logo and "cutting through complexity" are registered trademarks or trademarks of KPMG International.