topdown
CONSULTING

Taking EPM to new levels with Oracle Hyperion Data Relationship Management

TABLE OF CONTENTS

Taking EPM to new levels
with Oracle Hyperion Data Relationship Management

Topdown Consulting

Introduction

Organizations all over the world are realizing the advantages of analytical solutions. They are also integrating their operational data into their analytical systems to further its benefits. Master Data is critical in every one of these operational, analytical, and reporting solutions and must be coordinated across all of the systems. Many organizations have no strategy to address this issue.

This white paper will discuss briefly the need for processes identifying common definitions, Oracle's Hyperion solution for managing master data, specifically Oracle's Hyperion Data Relationship Management (DRM) application and how to leverage it in enterprise systems. To help illustrate this further, real world examples for different ways in which DRM is integrated into the enterprise solution will be provided.

Key Components

There are several key components to a successful implementation of any Master Data Management solution. Two of the most important are:

1.   A comprehensive Information Governance Process

2.   A Master Data Management toolset.

The goal of an Information Governance Process is to create a consensus on master data and its meanings. The tool required to implement the Master Data Management integration must be capable of implementing the results of the governance process so it is not only captured but also distributed to the enterprise systems. The tool should provide a repeatable process around the master data solution including validations and verifications for conformity with the business rules.

Information Governance Process

Every organization should create and maintain an alignment between the operational and the analytical business data in order to ensure a consistent set of information. Without it, many definitions of the same term occur and cause considerable misunderstandings within management. In order to properly communicate across all systems, an Information Governance Process should be set up to develop common hierarchies, relationships, business rules, terminology, formulas behind the terminology across all of the enterprise—i.e., the definition of Net Income and what it includes. This is a key part of the process in accomplishing a unified view of the organization.

An executive group—the Data Governance Council—should also be formed. This group will consist of representatives with the official capability to make final decisions about the master data. They will responsible for signing-off on any Master Data decisions.

Data Stewards need to be assigned from all of the systems and organizations involved in the enterprise solution. They will be responsible for identifying, discussing, coming to consensus on and proposing to the Data Governance Council the common objects, definitions, etc. Initially, the data stewards will need to meet to develop the standards and address existing issues. The proposed standards and resolutions would then be presented to the council for final decisions followed by the implementation.

The Information Governance process can be time consuming initially, but it will pay dividends by mitigating or minimizing miscommunication, conflicting reports, data validation issues, integration issues, and meetings between executive management trying to resolve these issues. Once the initial standards are established, existing issues are resolved, and the new process is in place, data stewards will meet regularly to discuss new items with its recommendations submitted to the council as needed. This ongoing effort is part of a continual cycle, not nearly as grueling as it is establishing the initial consensus.

The outlined process is often trimmed down when DRM is implemented in a focused smaller environment, especially when DRM will be a recipient of data from another system—e.g., when an ERP is the initiator of the metadata. In this case, it is often necessary to handle differences, conflicts and issues in the metadata that affect downstream systems. One example is addressing duplicates. In some systems, such as Essbase, it may be necessary to prefix different dimensions and add the unique member's names as a suffix to descriptions in order to make them unique. This is one method too often used to ensure uniqueness. Had this been addressed in a bigger picture governance process, the members and aliases would already be unique and void of any special characters that cause problems for downstream systems. These additional rules can often be built into the initiating process as business rules. Without this upfront enforcement, the downstream systems or tools such as DRM will need to be able to address any issues this may cause.

Technical Solution

A toolset that can implement the results of the Information Governance Process is critical. The tool needs to be able to:

• Provide a centralized, consistent, single source of truth for metadata

• Manage a wide range of hierarchies

• Enable business users to maintain the metadata

• Allow creating alternate hierarchies not available elsewhere in the organization

• Include an approval process and change management all in an easy to use interface.

DRM provides all of these capabilities and should be an integral part of any Master Data Management strategy.

Many of the metadata projects are sponsored by the Finance organization where there is a mix of business users and super-users and this is an excellent environment for DRM. Many of the super-users will be able to function in a development capacity building properties and validation and verification rules that are queries run to ensure that the metadata meets business rules. These may range from account numbers all having 10 numerical digits to more complex parent-child dependencies. They may also be involved in building out custom measures used in downstream systems that include formulas. The business users are relied upon to create or validate the new members, their proper locations in the hierarchies and set properties for the members. Security can limit what members and properties the business users can view, and which ones they can change.

Change management/control and audit/compliance are all included in the DRM toolset. All changes to any hierarchy, member or property are individually captured and are available in a searchable log. Different versions can be created from an existing version of the hierarchies from any point in time and new copies of the version may be created once a version is finalized in production and work on the next month begins. A status on the versions can be set to limit who, if anyone, can make changes to it.

Many more capabilities that are too numerous to mention exist in DRM. It is a capable tool and when integrated within a complete solution is very powerful.

Financial Data Repository (FDR)

As a part of many master data management solutions, relational databases can and often do play a very useful role. They can serve as a core tool in your Metadata Management process. It can provide a central point of access, maintenance, preparation, historical backups, comparisons, and ease in retrieval for all of your metadata. This central database integrating data and metadata is what TopDown Consulting refers to as a Financial Data Repository (FDR). Many companies already have a relational database in their enterprise solution, so the FDR leverages not only an existing core tool already in the Enterprise Architecture but an existing skill set as well.

One of the ways to leverage a relational database during the master data management process is to prepare the master data for use in DRM, though DRM can also do a lot of the tasks described, it is often better to help reduce its workload by completing some of these tasks prior to loading into DRM. If a company is using DRM, a relational database can reduce the amount of metadata that needs to be loaded into DRM by identifying only the changes that need to be brought in instead of bringing everything in, and then using DRM to identify the changes. It can also serve as the target from DRM and provide access to any applications capable of pulling their metadata from tables.

One method for identifying new master data members is to bring it into the database from the initiating system—e.g., an ERP—where comparisons may be made with the previous version, and differences can easily be identified often by using simple SQL commands. This is extremely useful when large metadata sets are involved. Stored procedures and views may be leveraged to build out files for loading into the downstream systems. These stored procedures also can easily provide DRM with just the changes including new members, moves, and changes to properties in a simple file that can be loaded. If DRM is not the master data management solution used by the company, the data can still be prepared to support whatever system will receive these changes/updates.

Another advantage in using a database is in retrieving the resulting hierarchies. A query can build the hierarchy in the correct order for loading into a target system. If there was a sequence when it was loaded, this can be maintained throughout the process. If there wasn't an order before, the hierarchy can be built recursively from the top down so it loads correctly into the target systems. This is a huge advantage for many systems! It also helps tremendously in maintaining and validating the hierarchy. Any changes can be verified for the correct depth of the nodes, shared members, correct connections from parent to child ensuring inclusion of all of the members, etc.

These are just some of the advantages to leveraging an existing asset and skill set as a component of the master data management solution. It is versatile and can help reduce the overall effort of the solution by preprocessing changes and in the variety of ways it can format the master data for subsequent use by downstream systems.

Real World Examples:

In the following real world examples, two different methods of utilizing DRM (formerly MDM in these diagrams) are shown. One way is to take a feed from the Enterprise system, import this into DRM, move the new members and changes into the proper location in the hierarchy, leverage DRM properties in the export and push downstream. The second is to initiate the Master Data members inside of DRM, assign properties and push to downstream. This is how the tool is really designed and leverages its strengths the best, though many organizations cannot go straight to this and have to take the hybrid approach of the two. Often there are established processes around an ERP system, for example, which exists and to move to using DRM as the initiator and build all of the interfaces required to retrofit these systems is unlikely, especially in the short term.

Large International Energy Services Company

In this example, most of the master data is created within SAP and then pushed into an Operation Data Store (ODS) within an Oracle database. Batch scripts were written to extract the metadata in a
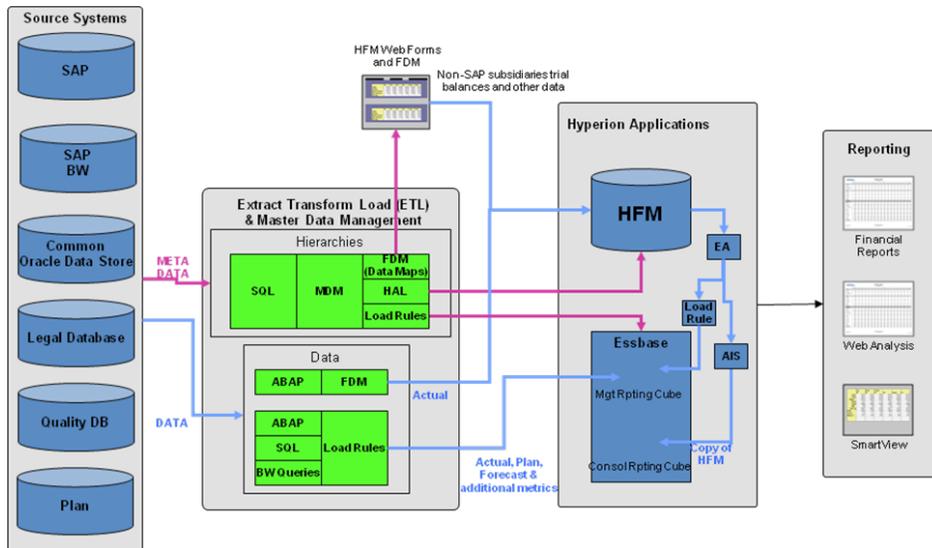
parent child format with aliases from the database into formatted input files for MDM (DRM). These files were then imported into MDM (DRM) and a blend was performed to add new members and move any existing members and update the aliases. The legal entity structure had to be pulled from a separate SQL Server system, which was done in a similar fashion.

The account dimension was handled in two different ways, one for an Essbase reporting cube and another for HFM (Oracle's Hyperion Financial Management software). Since the accounts from the ODS were all located within account sets, the Essbase account hierarchy was built so that the account sets

were included in the correct places within the hierarchy, though they were not included in the hierarchy export. This allowed the blend within DRM (MDM) to properly place the accounts under the account sets. If a new account set came in, it was identified within a different comparison. This batch process to identify new members was built so that it compared the new members with the blended hierarchy and exported differences into a DRM (MDM) import file. This file was subsequently imported back in and the new accounts were now located under a special member used to store new members. From this location, a user can drag and drop the member into the correct location with all of its correct properties. In addition, there were many custom measures included in the reporting cube's account hierarchy that did not apply to HFM. DRM can be used to store not only the members but also the formulas as well.

The HFM account hierarchy had a few different requirements. All of the cash accounts, where brought together under just a couple of accounts. This was accomplished by combining an export from DRM (MDM) that mapped all of these cash accounts to these new parents within FDM. FDM allows for these explicit mapping files to be loaded. These are csv files containing the 'from account' and 'to account' with possibly another property depending on the dimension. These base cash accounts were subsequently not pushed from DRM (MDM) into HFM as only the parents were now needed. In addition, the hierarchy itself was very different, so a similar process to the one described above identifying the new accounts was used to prepare accounts for placement in the HFM hierarchy. Since the siblings of these accounts where known in the initial load of the new account hierarchy from SAP, the DRM (MDM) administrator could propose the locations in HFM which were often correct unless it was identified as something that needed to be tracked separately, for example.

Large Oilfield Services Company

In this example, JD Edwards is the system of record for many of the hierarchies. Instead of an
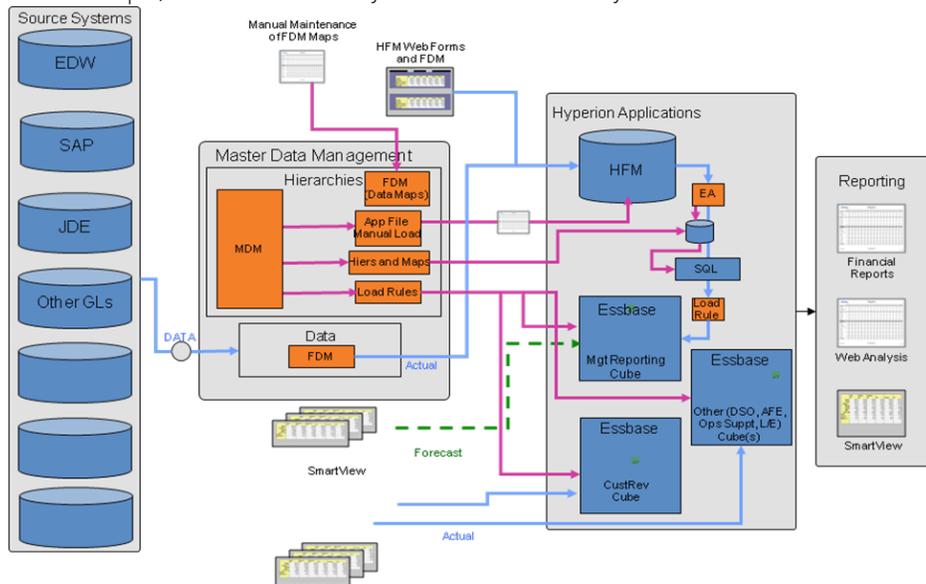
automated pull, though, the same team that manages the JD Edwards hierarchies also manages DRM and inputs the new members in both places at the same time either manually, but more often by using the Automator feature. This feature allows you to "import" a csv file with the members, their parents and properties defined. The "import" actually makes the calls required to the API interface to execute the actions you define for it in the automator file. This may include adding a node, updating properties, deleting or moving a node, and many more options. This is a very powerful feature which can be easily combined with a relational database to create a simple solution used to update master data from external sources. This was discussed some in the FDR section above. If the master data is loaded into the database, it can be compared to the previous load and an automator file can be built containing the differences. This file is then run into the new month's version of the hierarchies. This can be written into a batch job so that all of the changes are made without manual intervention required.

There were actually two targets for the master data in this architecture. The first was HFM, which requires a file called an APP file be built and pulled into HFM containing all of the members, their relationships and properties. This was accomplished using the "Book" export feature which combines many exports into one file including a header and footer. Exports for the individual sections were created and combined in the correct order into this Book which can be exported with one command or click. The other target was a set of tables in an oracle database which, when combined with output from the Extended Analytics (EA) feature of HFM was used to create the Essbase Reporting Cube. Extended Analytics (EA) is an export capability within HFM which can create a set of tables in a star-schema format. Using the web interface, a user can select data to be extracted from various intersections within HFM to the set of tables. They also include the metadata that can be leveraged to create a cube as well. In this case, there were several alternate hierarchies and rollups that were included in the reporting cube but not in HFM. Pushing these hierarchies directly to oracle tables for use in creating the Essbase cube provided the opportunity to create these without cluttering up HFM. Custom measures also were included in the reporting cube's account hierarchy which did not apply to HFM. The Book for HFM and the output to these tables where done at the same time to ensure integrity.

Summary

An MDM strategy that includes an Information Governance Process in conjunction with Oracle's Hyperion Data Relationship Management (DRM) eliminates the siloed approach to preparing master data and measures for reporting/analytics and transactional systems. It forces a consistent set of business processes that will guarantee master data integrity and uniformity. DRM provides the capability to build in these business rules both for when you are creating the master data within the tool as well as when you are loading it and need to check it afterwards. By implementing this strategy, the organization will reduce the time spent preparing the master data for use by the BI reporting environments. More important, it will allow Business Analyst to get out of the master data integration business and back to doing the business analysis on a quality analytical solution.

About TopDown Consulting

About TopDown Consulting, Inc. Founded in 2000, TopDown Consulting is the acknowledged leader in designing, implementing, and deploying EPM solutions. TopDown has the experience, expertise, and proven approach to deliver successful implementations for Global 2000 clients. Our consultants average 20 years of Hyperion and industry experience with a complete range of skill sets, including: solution and system architects, CPAs, finance executives, MBAs, analysts, and project managers. TopDown's proven Project Success Methodology provides a customizable strategic framework for guiding and measuring project initiatives, enabling us to deliver solutions that meet our clients' current and future business and technology needs and their unique corporate culture.

TopDown Consulting, Inc. serves clients nationally and internationally from our San Francisco headquarters.
For more information or to inquire about our services, please contact us.

TOPDOWN CONSULTING, INC.

530 Divisadero Street #310   Phone: (888) 644-8445 (Calls from within US and Canada)   Email: info@topdownconsulting.com
San Francisco, CA 941177   Phone: (617) 576-8445 (Outside US and Canada)   Web: www.topdownconsulting.com