


Oracle Warehouse Builder 10g and OLAP – What's New

Oracle OPENWORLD '05

Session #675



Chris Claterbos

claterbos@vlamis.com

Vlamis Software Solutions, Inc.

<http://www.vlamis.com>


Copyright © 2005, Vlamis Software Solutions, Inc.



Vlami Software Solutions, Inc.

- **Founded in 1992 in Kansas City, Missouri**
- **A Member of Oracle Partner Program since 1995 along with various Oracle Beta Programs**
- **Designs and implements databases/data marts/data warehouses using RDBMS and Multidimensional tools**
- **Specializes in Data Transformation, Data Warehousing, Business Intelligence, Oracle Financials and Applications Development**
- **Founder Dan Vlami is former developer at Oracle-Waltham office for Sales Analyzer Application**
- **Oracle Certified Solutions Provider**





Who Are We?

- **Chris Claterbos, Development Manager**
 - ❑ DBA and applications developer for Oracle products, since 1981.
 - ❑ Beta tester and early adopter of - including Oracle 8i, 9i and 10g, JDeveloper and BIBeans, Oracle AS, Portal (formerly WebDB), and Reports.
 - ❑ Speaker and author.
 - ❑ Previous IOUG Focus Area Manager for Data Warehousing and BI
 - ❑ Consulting and Development Manager for Vlamis Software Solutions, Inc.



Using OWB to Create OLAP Databases

- Introduction
- Oracle 10g and OLAP
- What is OWB?
- What is New in Paris?
- Oracle 10g Integration
- Design objects
- The Process
- Demonstration
- Managing an OLAP project
- Getting Started
- Questions



2005 is the YEAR of BI - OLAP



But not for OWB Paris!

Sorry!

We Do Have a PLAN!

- [About OTN](#)
- [Oracle ACEs](#)
- [TechBlast Newsletter](#)
- [Oracle Magazine](#)
- [Blogs & Opinion](#)

updated August 24, 2005

Learn the advantages of adding the Oracle Database Enterprise Edition OLAP Option to your Oracle BI Discoverer reporting and analysis, and the steps required to make OLAP

- [Oracle Business Intelligence Technical Overview \(PDF\)](#)
- [Oracle Business Intelligence Demo](#)
- [Oracle by Example: Online BI Tutorials](#)
- [Data Warehousing & BI Discussion Forums](#)

- x [Oracle Business Intelligence 10g](#)
- x [Analytic Workspace Manager](#)
- x [Oracle Warehouse Builder](#)
- x [Oracle Data Miner](#)
- x [Oracle Database 10g Enterprise Edition](#)

- . Query & Reporting
- . Data Warehousing Fundamentals
- . Introduction BI Solution
- . Reports Development
- . BI Developer Track

Internet Seminars

Customers need a Unified View

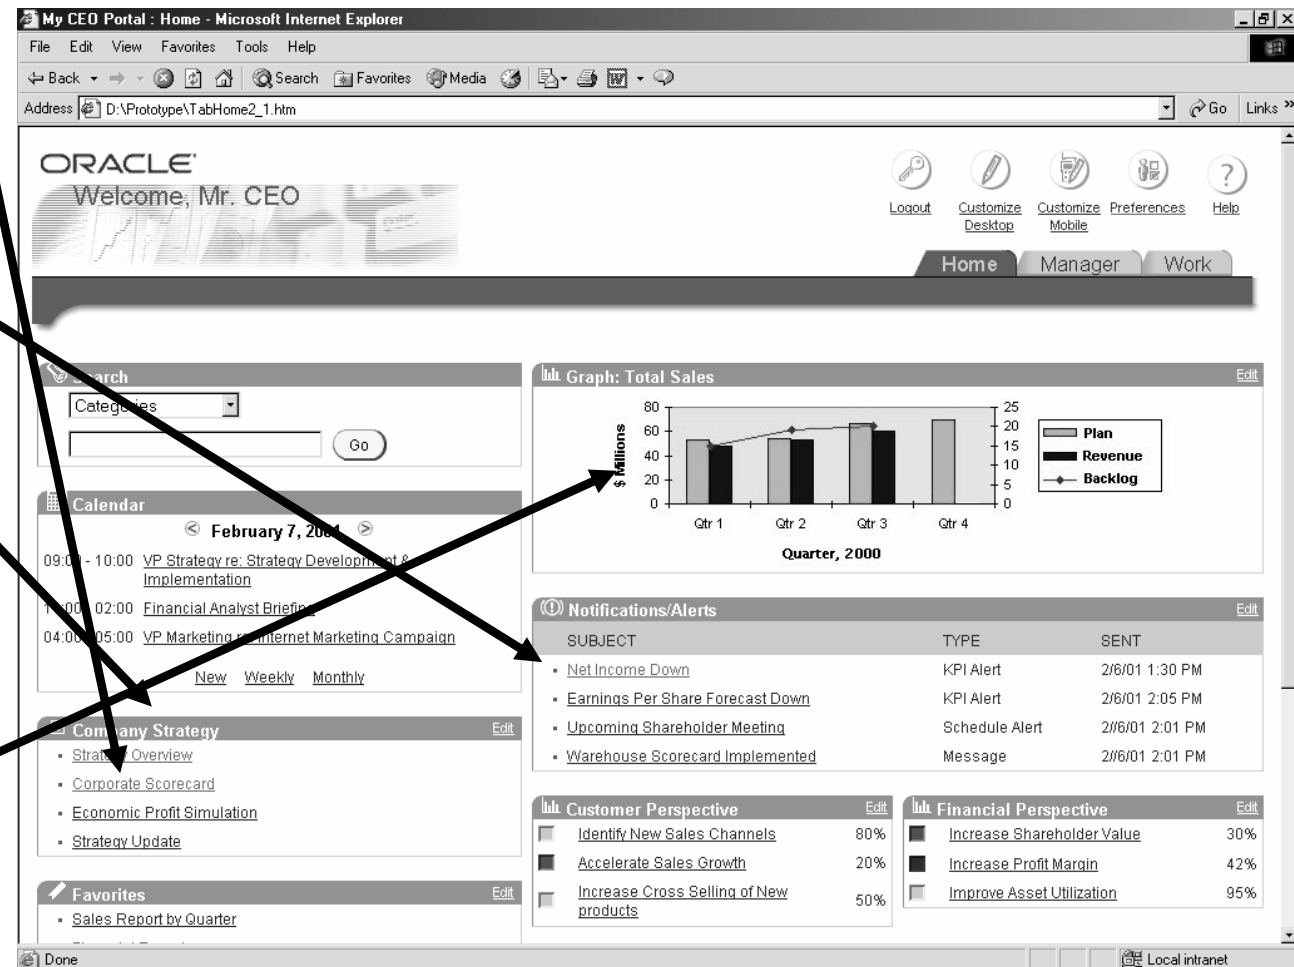


Planning

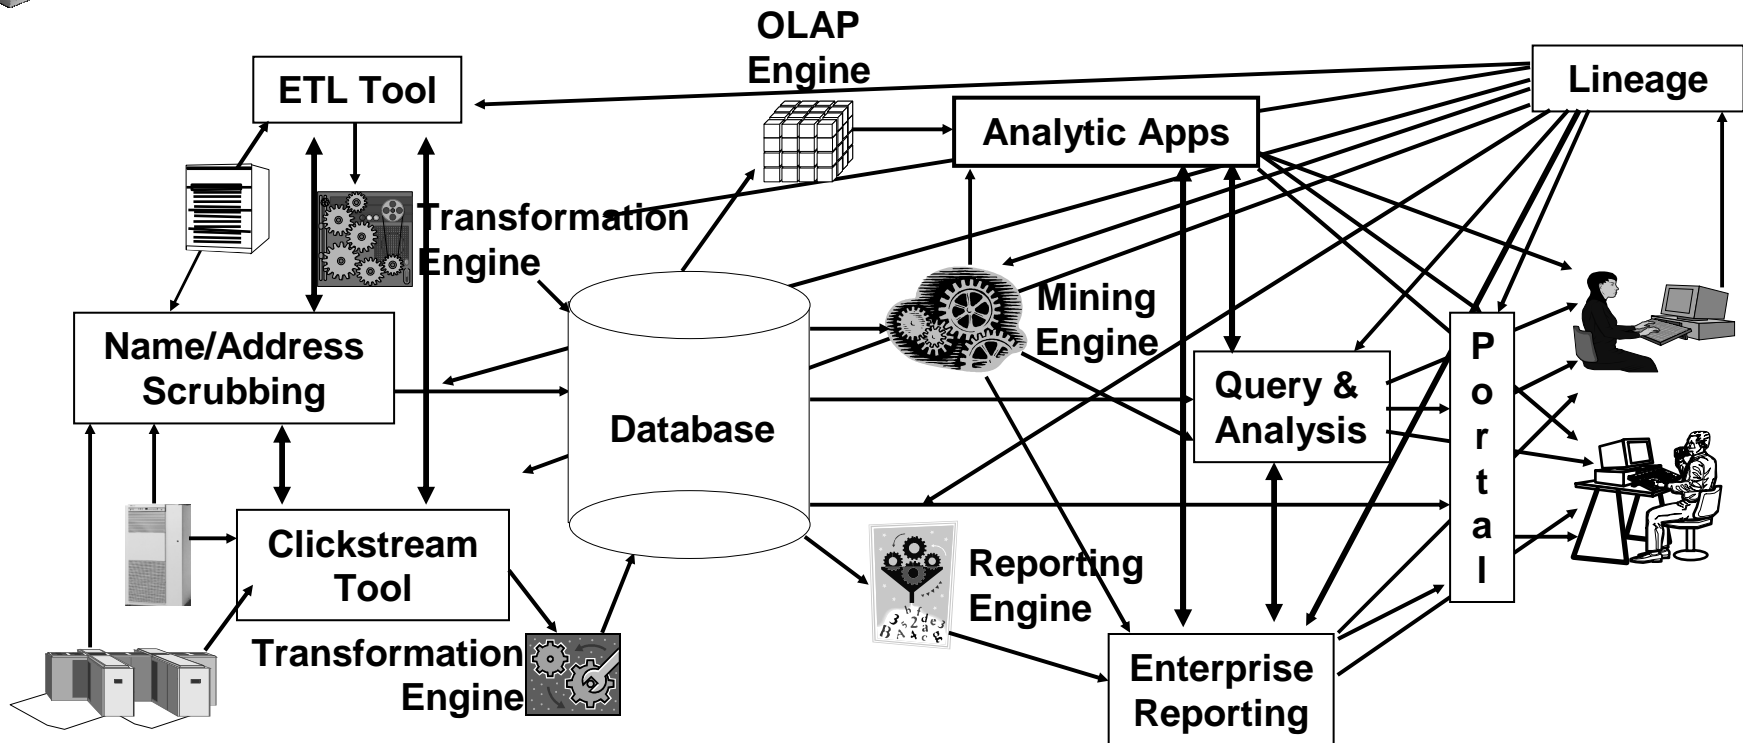
Monitoring

Analysis

Reporting

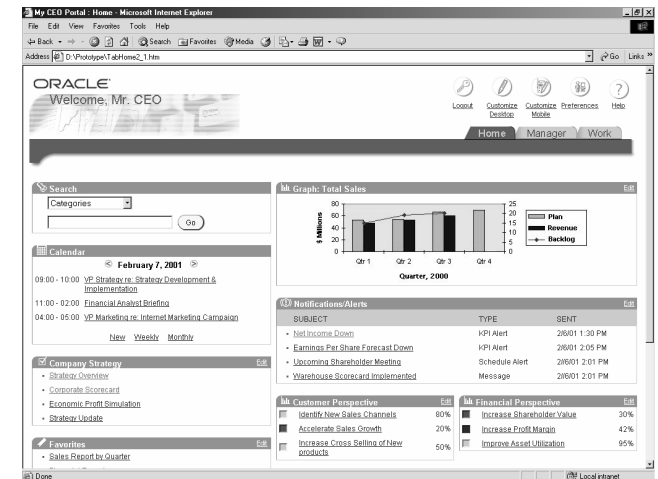



Business Intelligence Market Multi-Vendor, Un-integrated




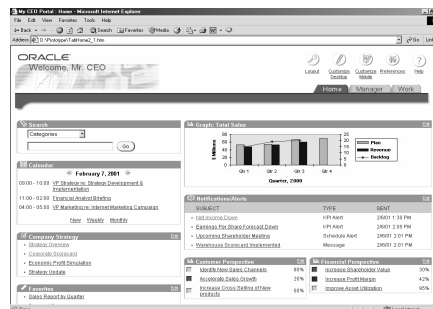
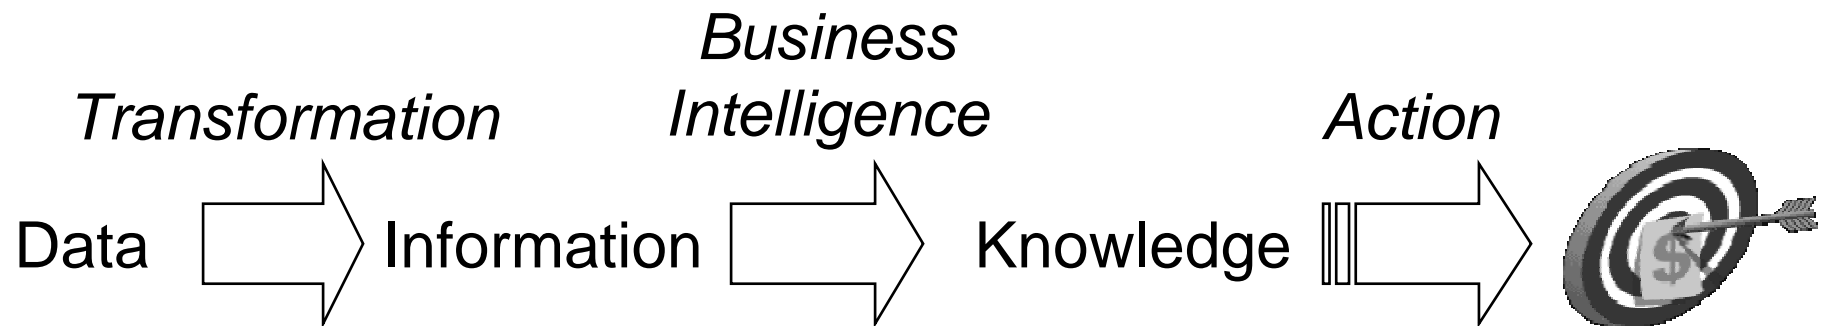
- Protracted and complex implementation
- Escalating maintenance costs
- Software ***and Metadata*** Integration is key!

How do I get from Raw Data to a Unified View?



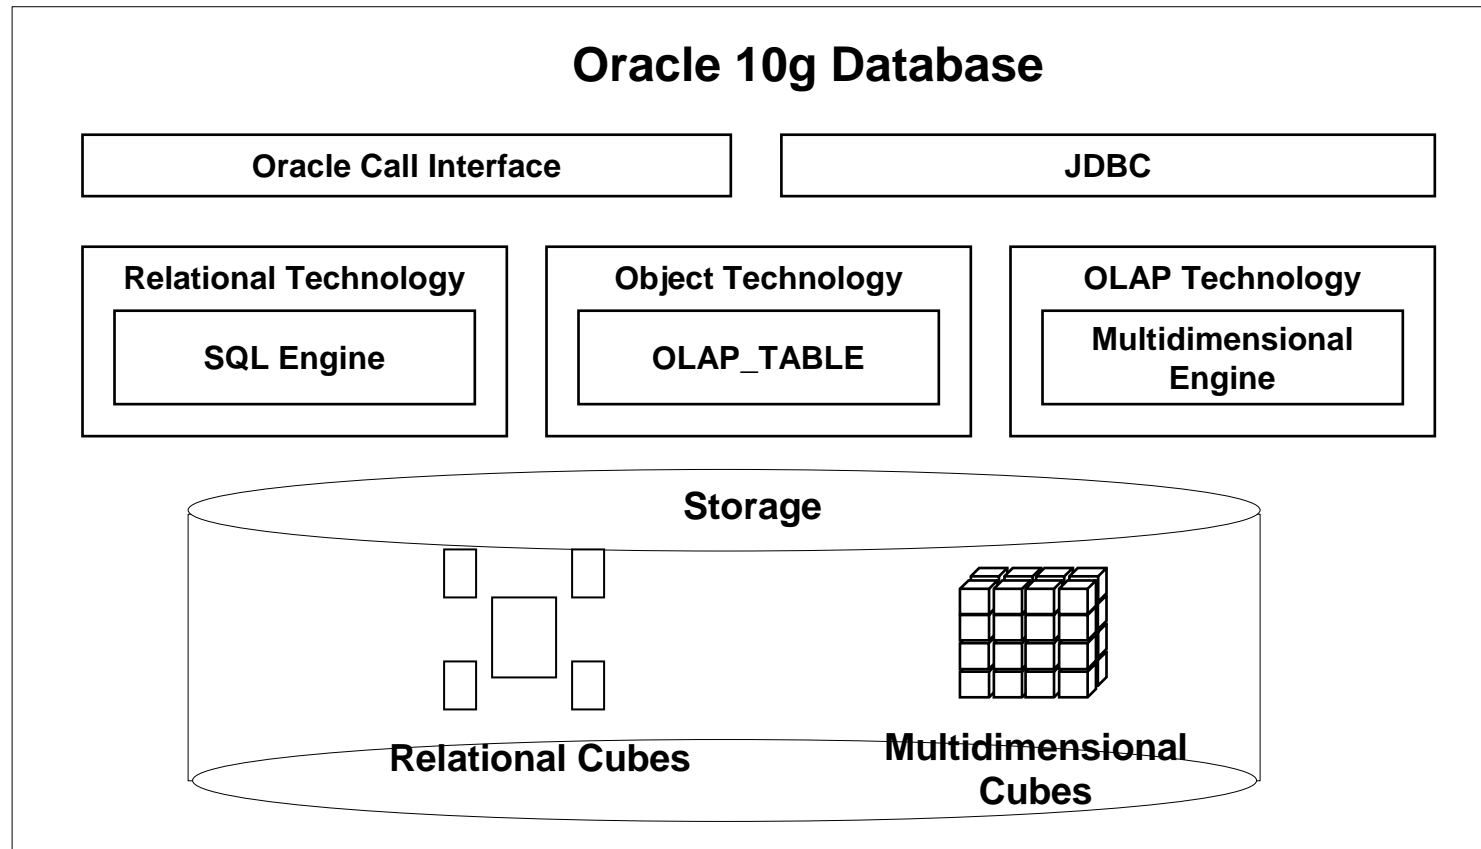


Turning Data into Profit....





Oracle 10g RDBMS - MDDS






What Does 10g OLAP Add?

- **Multidimensional user view of data**
- **Users create own reports**
- **Users create own measures**
- **Easy drill-down, rotate**
- **Iterative discovery process (not just reports)**
- **Ad-hoc analysis**
- **Easy selection of data with business terms**
- **What-if, forecasting**



What Makes a DW OLAP-Ready?

- **Star or Snowflake schema design**
- **Simple or complex dimension tables (level-based)**
- **Each child has single parent (no many-to-many)**
- **Total level at top of each dimension (except Time?)**
- **End_date and Timespan attributes for TIME**
- **Unique descriptions across all levels**
- **Fact tables with additive measures**



Why OWB to build OLAP?

- **Integrated with entire Oracle stack**
- **Graphically designs, generates, and deploys**
- **Only ETL tool that understands Oracle OLAP**
- **Uses 10g PL/SQL for transformations**
- **One-click deployment of 10g OLAP AW**




What is Oracle Warehouse Builder?

- **Integrated Tool for Data Warehousing**
- **Based on Common Warehouse Metadata Standard (OMG)**
- **Supports Design and ETL Functions**
- **Enterprise Framework for Designing and Deploying Datawarehouses and Datamarts**
- **Future integration platform for Express**



Key OWB Paris themes

- **Improved User Interface**
- **Enabling Quality Information**
- **Enabling Business Intelligence**
- **Enabling Expertise capture**
- **Signification improvements in usability and functions over previous releases**



Sources & Targets

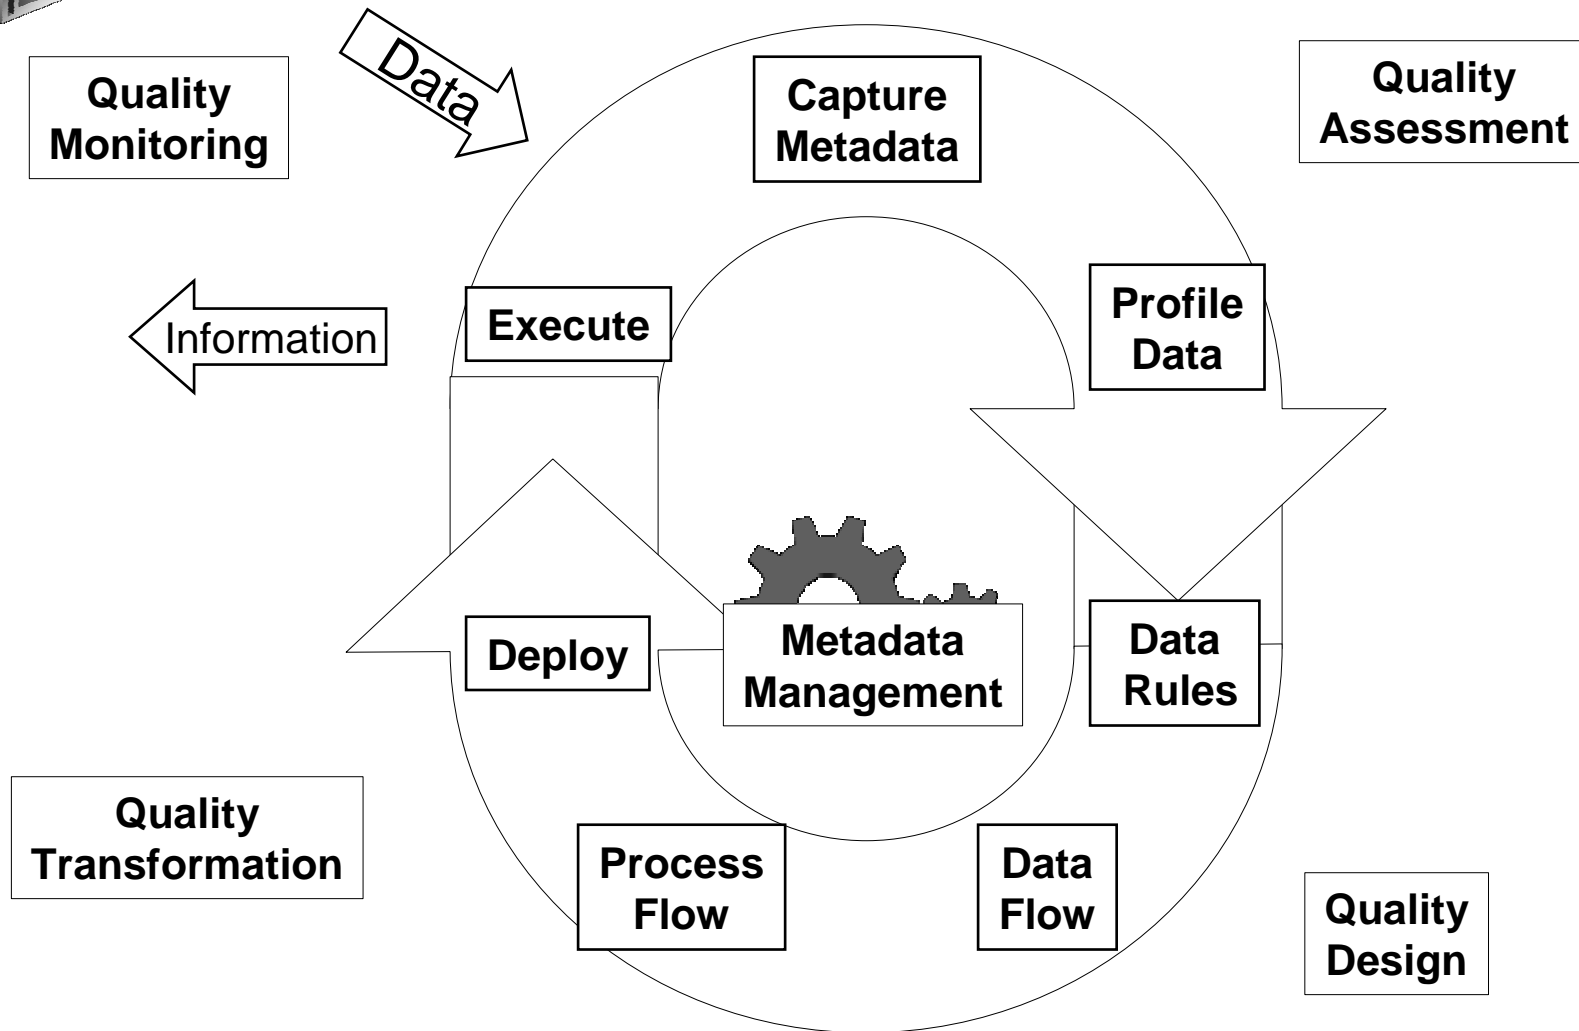
Sources

- Oracle
 - ❑ Tables, Views, MViews, Queues, External Tables, Table Functions, Streams, PL/SQL API's, Sqlloader...
- DB2, Sybase, SQLServer, Informix, ... (Oracle Transparent Gateways)
- Any ODBC source
- Flat Files
- Applications
 - ❑ Oracle Apps
 - ❑ SAP
 - ❑ Custom SQL App

Targets

- Oracle
 - ❑ Tables, Streams, OLAP, Table Functions, PL/SQL API's
- DB2, Sybase, SQLServer, Informix, ... (Oracle Transparent Gateways)
- Flat files

Enabling Information Quality



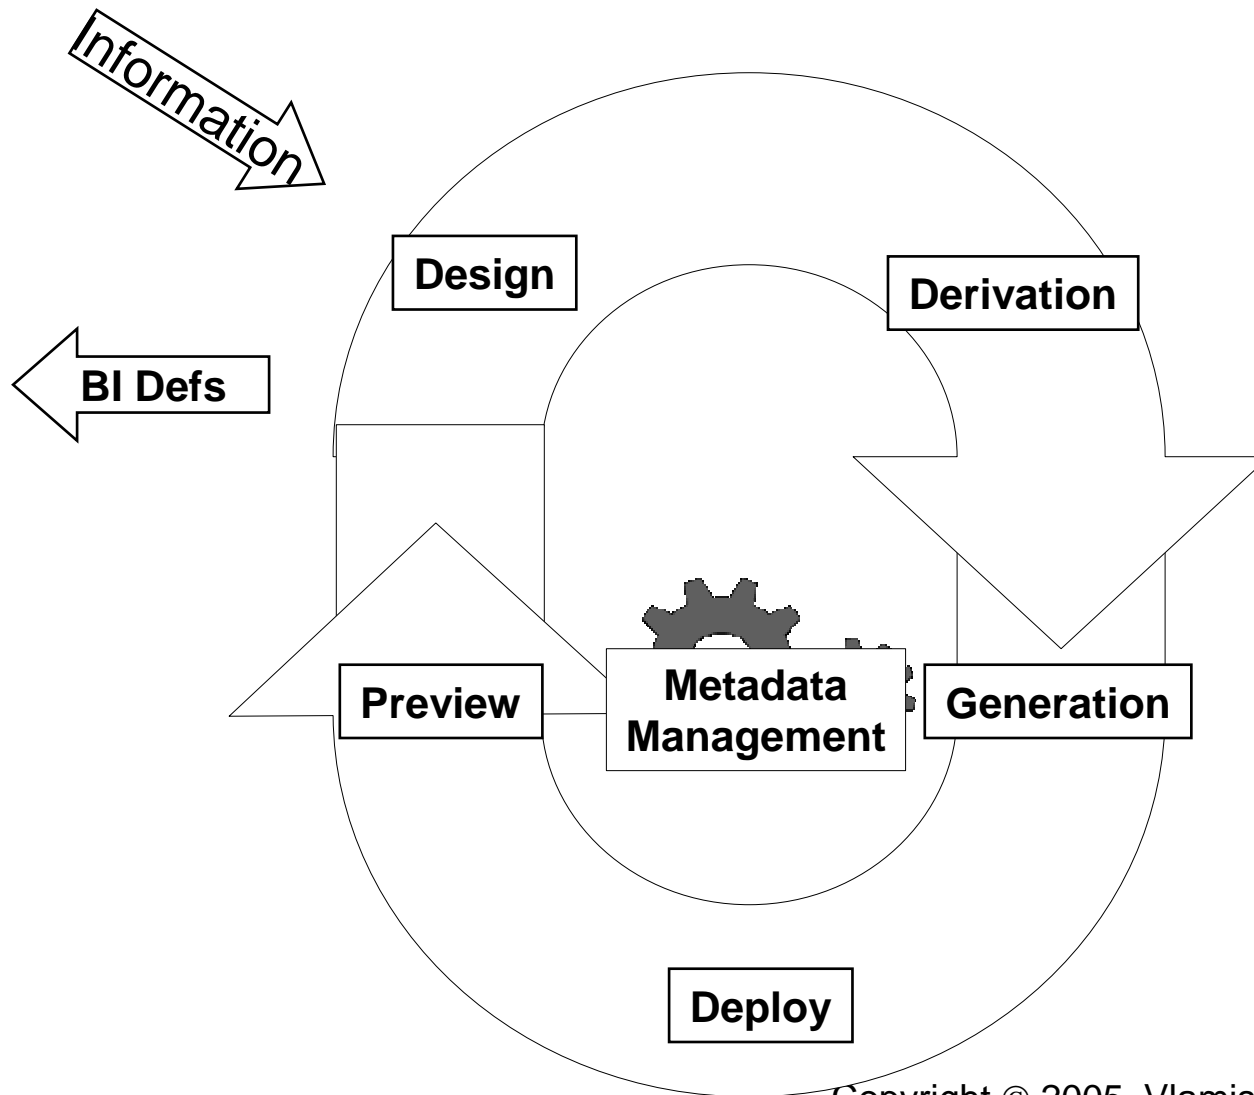
Data Object Design – One Editor



- Dimensions, cubes, tables, views, complex objects, ...
- Support for Star, Snowflake, Skip-Level, calculated measures, ...
- One editor for creation, configuration, validation, code generation, impact analysis, deployment, data viewing



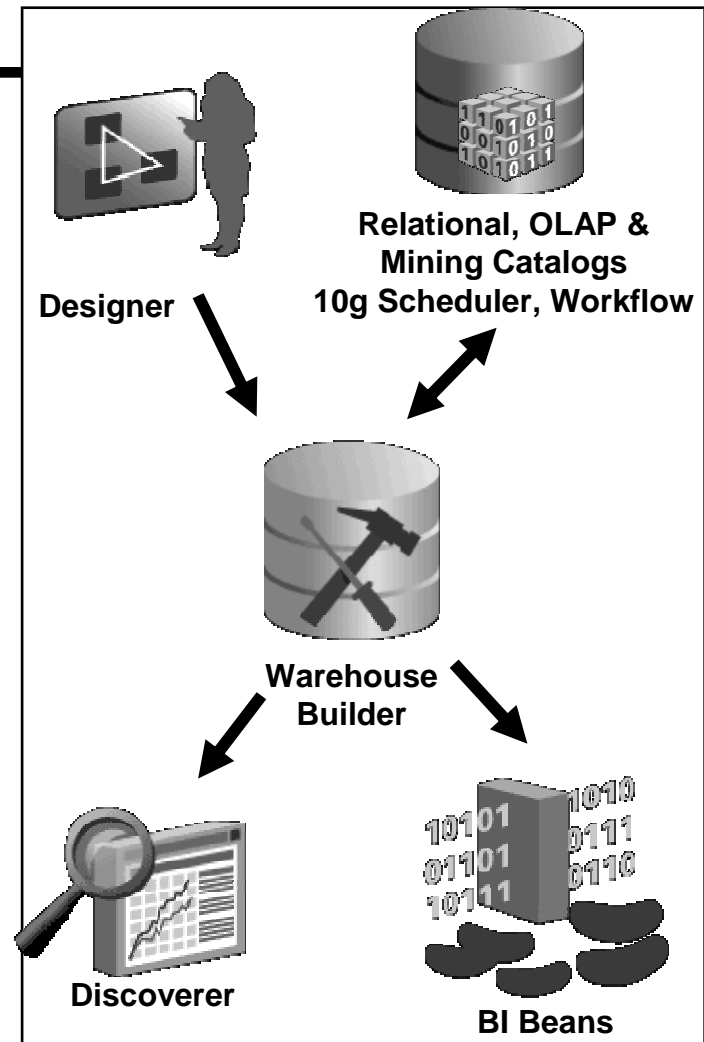
Enabling Business Intelligence



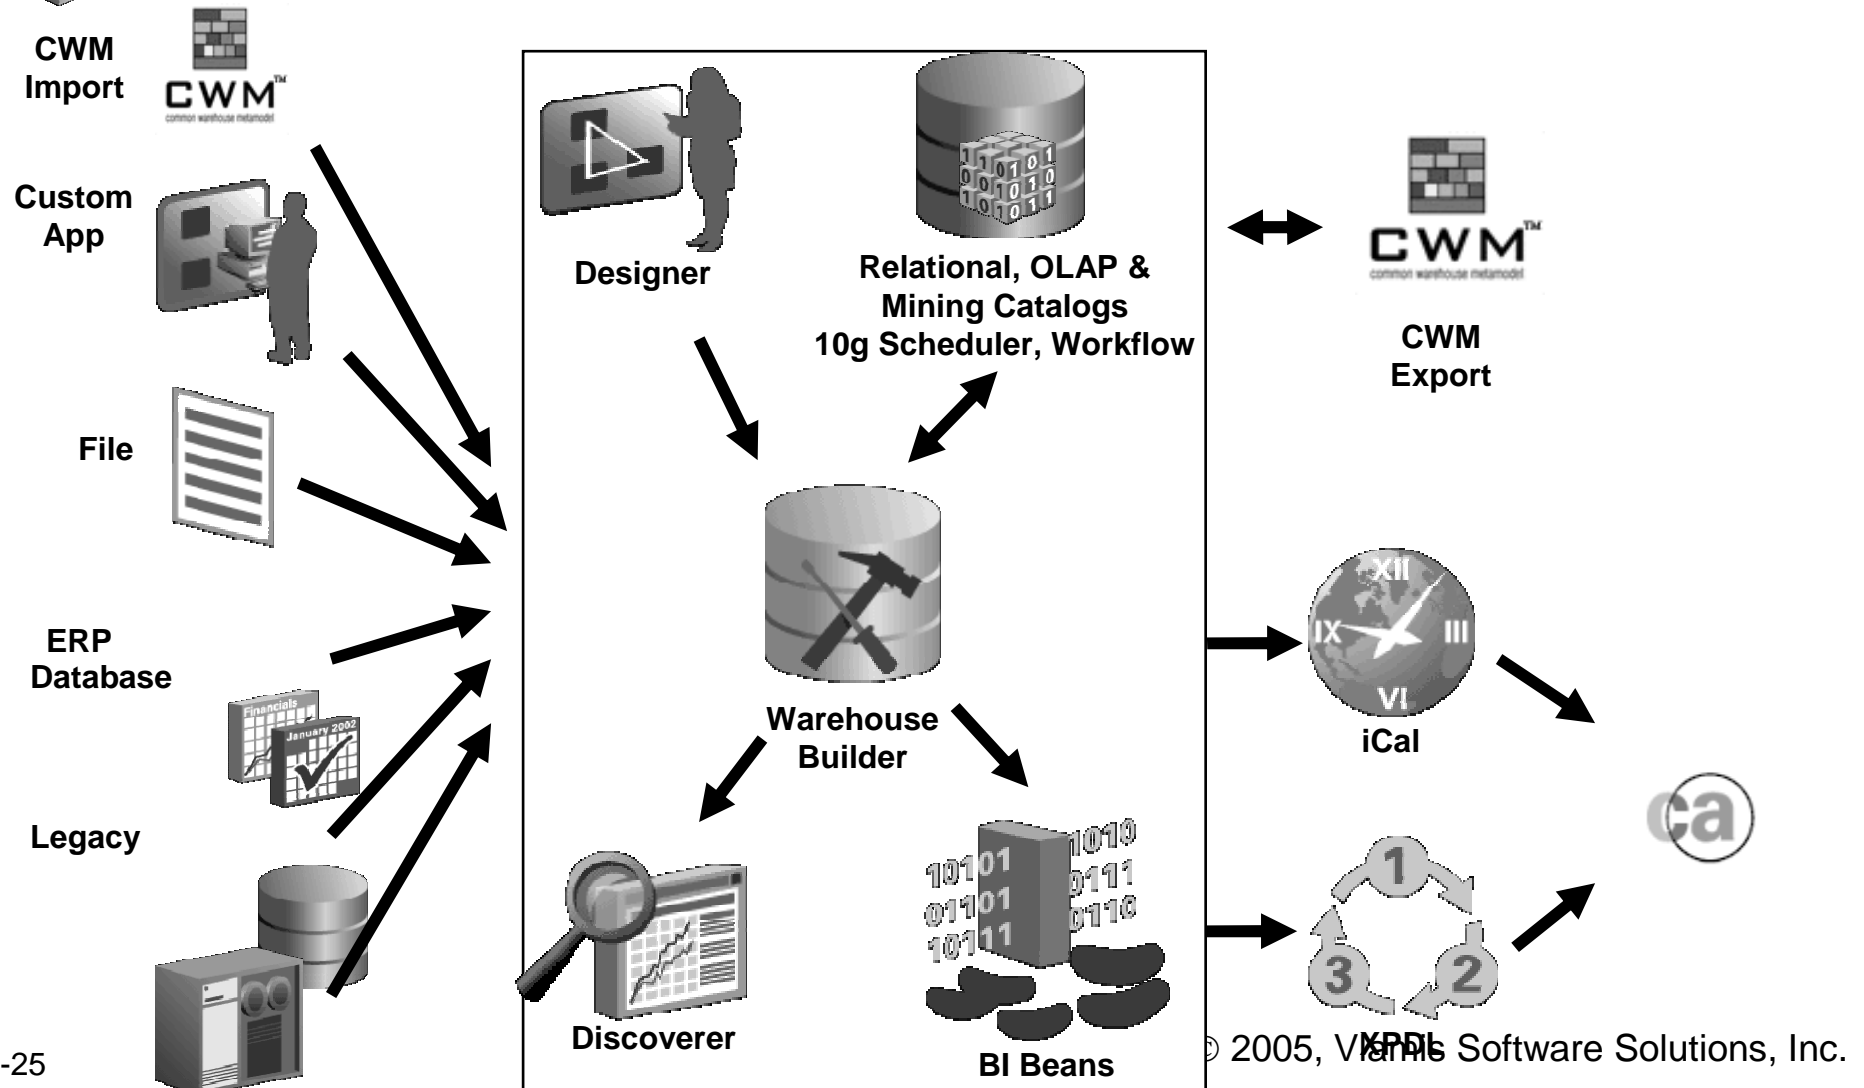

Business Intelligence Object Derivation



- Create and Derive Business intelligence objects
 - Oracle OLAP Cubes & Dimensions
 - OracleBI Discoverer EUL
 - OracleBI Beans Reports
- Included in Lineage and Impact analysis!



End-to-End Meta Data Integration





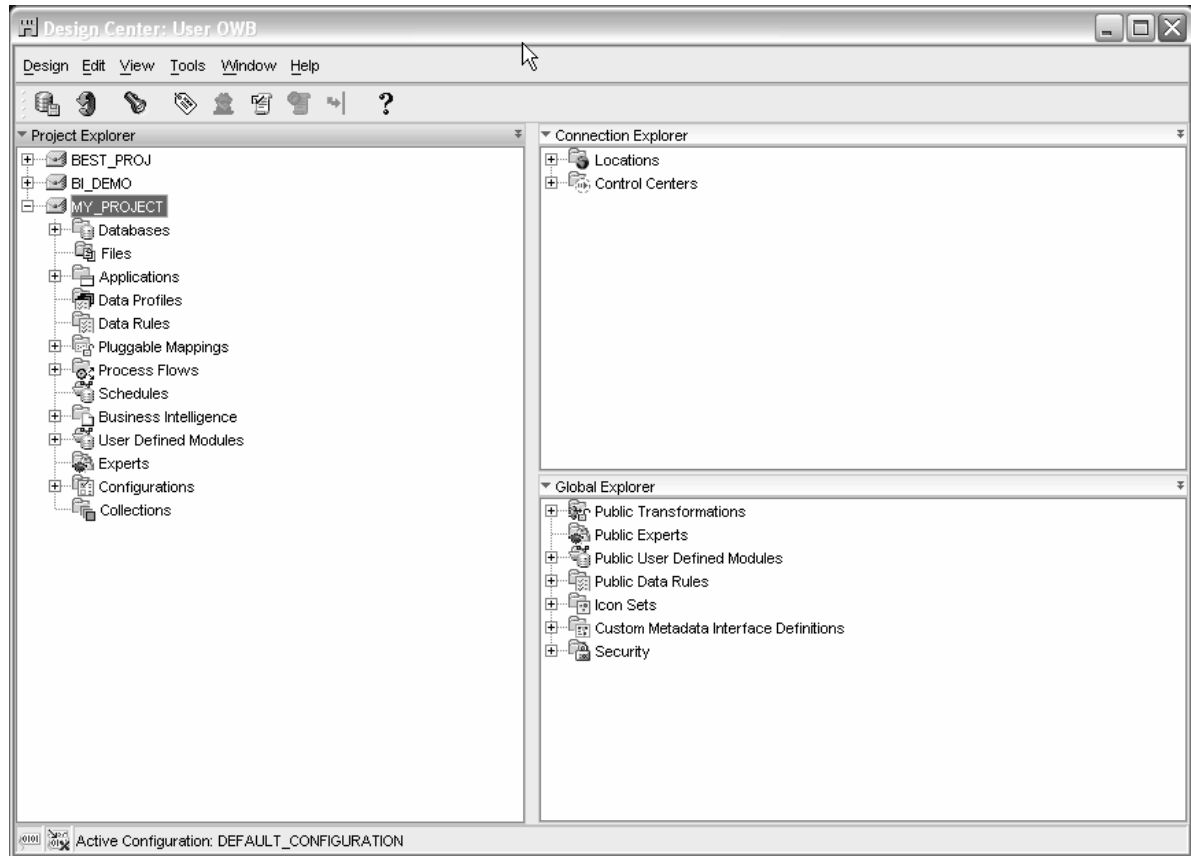
Platforms, Packaging

- **Available On:**
 - ☐ Win32(Windows NT/2000/XP/2003), Win64(XP/2003), Linux x86, Linux Itanium, Solaris, HP-UX (RISC), HP-UX (Itanium), AIX, Tru64
- **Packaging:**
 - ☐ Oracle Developer Suite (iDS)
 - ☐ Oracle Business Intelligence
- **Release date CY 2006**



Components: *OWB User Interface*

Java Based
Same look and
feel as Designer
Significantly
simplified over
previous
versions





Components: *OLAP Wizards*

- **Full Integration support for OLAP AWs**
- **Supports 10g OLAP as a Target**
- **Full Life Cycle support**
- **Viewing data is integrated into new interface**



10g OLAP Integration

- **OWB metadata to Oracle OLAP Metadata**
- **Create ROLAP or MOLAP objects**
- **Creates links to Relational Data for Facts and Dimensions (views or tables)**
- **Creates Scripts for building Materialized Views that are BI Beans OLAP friendly (ROLAP only)**
- **Creates Scripts to build and populate Analytic Workspaces**
- **User can use AWM to make changes (but cannot reverse engineer)**

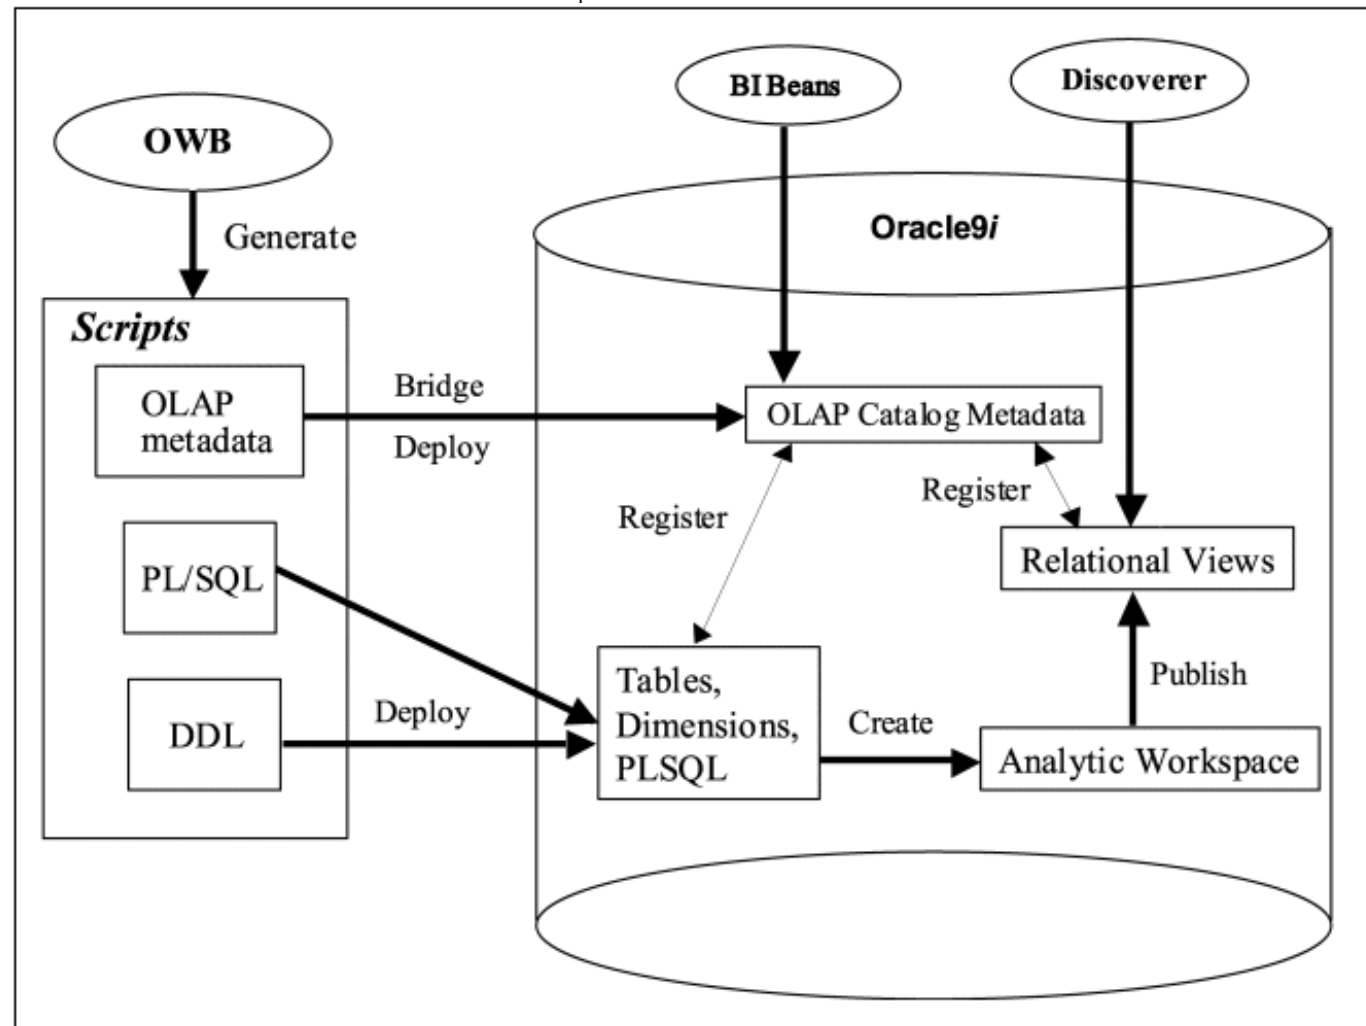


10g R1 – Uses Bridge to Integrate

- **OWB metadata to Oracle OLAP Metadata**
- **Create ROLAP or MOLAP objects**
- **Creates links to Relational Data for Facts and Dimensions (views or tables)**
- **Creates Scripts for building Materialized Views that are BI Beans OLAP friendly (ROLAP only)**
- **Use Bridge to Build AW Cubes and Dimensions**
- **Or use AWM to Map to Tables (BEST WAY)**
- **Used Pre-Defined Process to populate AW**



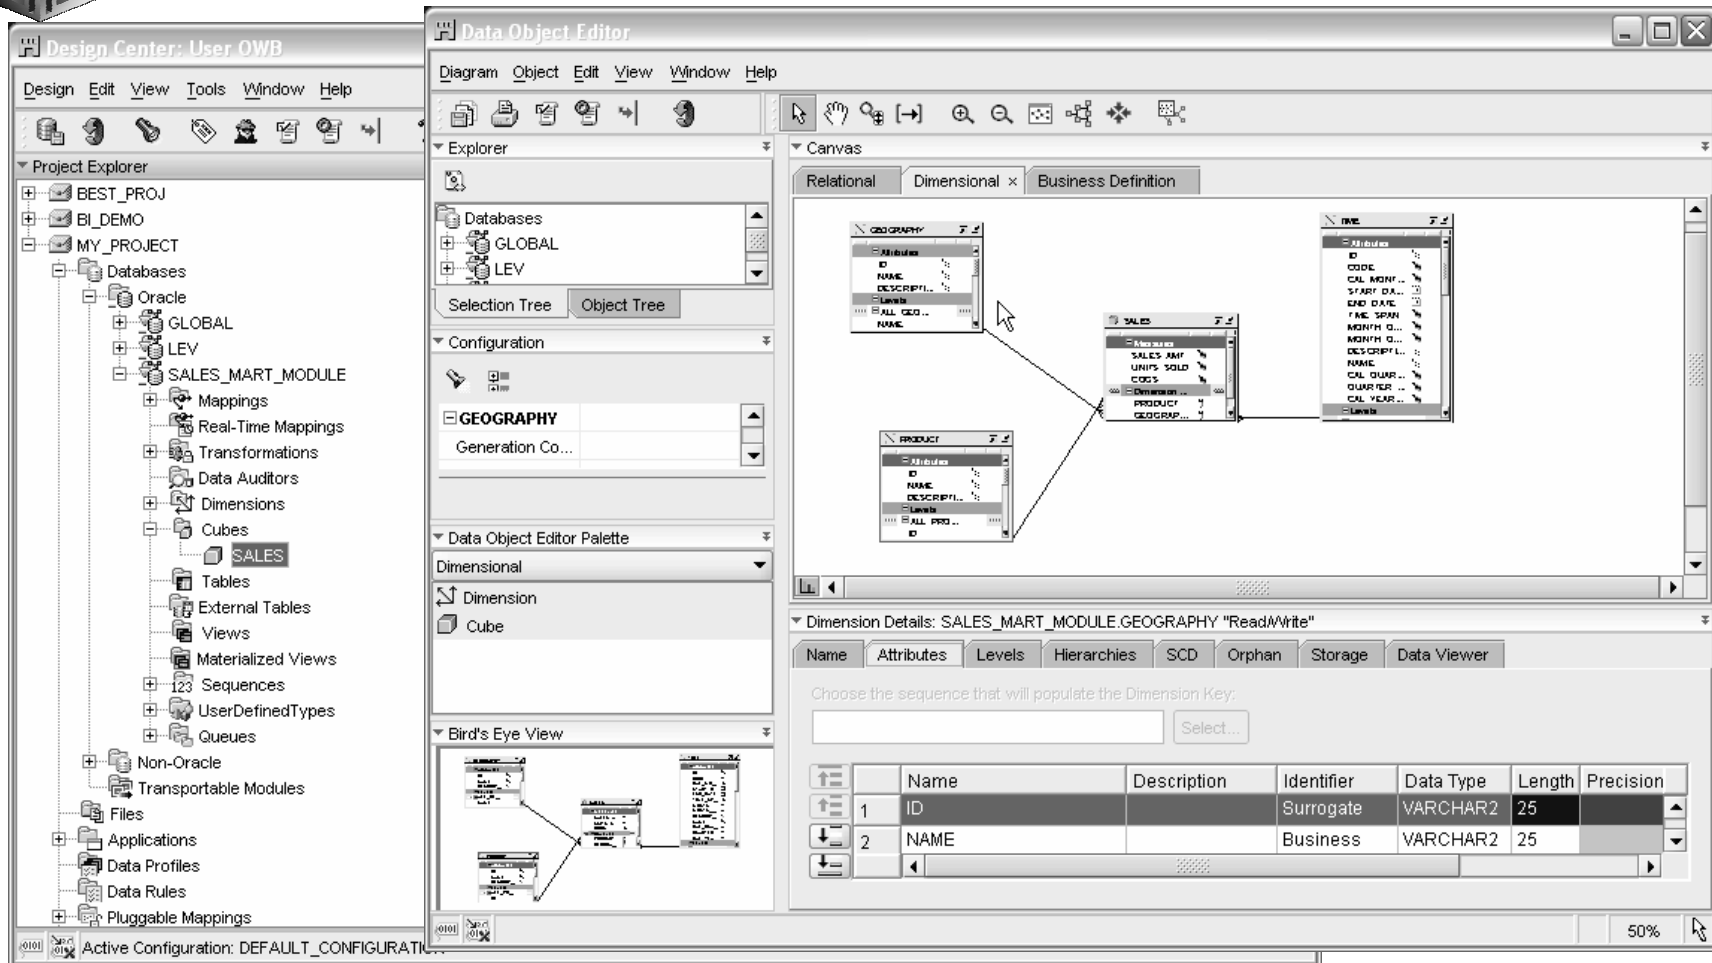
OWB OLAP Bridge





The Process

- **Design or Import Relational Schema**
 - ☐ **Define Fact Table(s)**
 - ☐ **Define Dimensions**
 - ☐ **Define Cubes (collection of like measures)**
- **Create Physical Schema**
- **Create Script for 10g OLAP or JUST DEPLOY!**
- **View/Modify in OWB**
- **Run Application**
- **Gather Statistics / Tune**





Creating Dimensions

- Use OWB to Create Dimensions
- Use the following “Special” Attributes when building OLAP Dimensions

Physical Level Attribute Name Suffixes in Warehouse Builder	Dimension Attribute Created
_NAME or NAME	Short_Description or Long_Description
_END_DATE or END_DATE	End_Date
_TIME_SPAN or TIME_SPAN	Time_Span
_PRIOR_PERIOD or PRIOR_PERIOD	Prior_Period
_YEAR_AGO_PERIOD or YEAR_AGO_PERIOD	Year_Ago_Period



Creating Dimensions

Table Properties: OWB_TIMEDIM_DATA_TABLE [Read/Write]

Name Columns Constraints Attribute Sets User Defined Properties

Table Columns

Name	Position	Data Type	Length	Precisi...	Scale	Not Null	Note
WEEK_OF_YEAR	15	NUMBER		0	0	<input type="checkbox"/>	
WEEK_START_DA...	16	DATE				<input type="checkbox"/>	
WEEK_END_DATE	17	DATE				<input type="checkbox"/>	
WEEK_TIME_SPAN	18	NUMBER		0	0	<input type="checkbox"/>	
MONTH_ID	19	NUMBER		0	0	<input type="checkbox"/>	
MONTH_OF_QUA...	20	NUMBER		0	0	<input type="checkbox"/>	
MONTH_OF_YEAR	21	NUMBER		0	0	<input type="checkbox"/>	
MONTH_START_D...	22	DATE				<input type="checkbox"/>	
MONTH_END_DATE	23	DATE				<input type="checkbox"/>	
MONTH_TIME_SPAN	24	NUMBER		0	0	<input type="checkbox"/>	
QUARTER_ID	25	NUMBER		0	0	<input type="checkbox"/>	
QUARTER_OF_YE...	26	NUMBER		0	0	<input type="checkbox"/>	

Add Remove

Help OK Cancel



Creating Time Dimensions

- Time Dimensions are “Special” Dimensions that allow for several analytic analyses such as “Sales last month compared with same month last year”
- Requires special attributes
- OWB has sample definition and SQL scripts for “Best Practice”
- Always use “Time” or “_Time” in Dimension Name – Like “T_TIME” or “TIME” (Paris only)



Creating Time Dimension

New Wizard to Create!

Time Dimension Attributes:

Physical Level Attribute Name Suffixes in Warehouse Builder	Dimension Attribute Created
_YEAR	Year Level
_QUARTER	Quarter Level
_MONTH	Month Level
_DAY	Day Level

Note: Week is not included because week cannot neatly rollup into calendar year.



Creating Dimensions

OWB now Supports Slowly Changing Dimensions!

- **Type 1 – Do not save history (default)**
- **Type 2 – Save History**
- **Type 3 – Store only previous value**
- **Supported by 10g OLAP!**



Creating Dimensions

OWB now Supports Ragged and Skip Levels!

- **Must Load Dimension via Snowflake (now default)**
- **Can have Ragged and Skip in same Dim**
- **MUST use 10.1.0.4 Target to work – 10.2 preferred!**

Can be done NOW with AWM 10.2!



Time Dimension

Data Object Editor

Diagram Object Edit View Window Help

Explorer

- Databases
 - GLOBAL
 - LEV
 - SALES_MART_MODULE

Selection Tree Object Tree

Configuration

- TIME**
 - Generation Comme...
 - Identification**
 - Deployable ☒

Data Object Editor Palette

Dimensional

- Dimension
- Cube

Bird's Eye View

SALES

- Measures
 - SALES AMT
 - UNITS SOLD
 - COST
- Dimensions
 - PRODUCT
 - GEOGRA...

TIME

- Attributes
 - ID
 - CODE
 - CAL_MON...
 - START_DA...
 - END_DATE
 - TIME_SPAN

Canvas

Time Dimension Details: SALES_MART_MODULE.TIME "Read/Write"

Name Storage Attributes Levels Hierarchies


Choose the sequence that will populate the Dimension Key:

TIME_SEQ Select...

Dimension Attributes

	Name	Description	Identifier	Data Type	Length	Precision	Scale	Second...	Descriptor
1	ID		Surrogate	VARCHAR2	25				
2	CODE		Business	NUMBER		0	0		
3	CAL_MONTH_NUMBER			NUMBER		0	0		
4	START_DATE			DATE					
5	END_DATE			DATE					
6	TIME_SPAN			NUMBER		0	0		
7	MONTH_OF_QUARTER			NUMBER		0	0		
8	MONTH_OF_YEAR			NUMBER		0	0		
9	DESCRIPTION			VARCHAR2	2000				Long descripti
10	NAME			VARCHAR2	25				Short descripti
11	CAL_QUARTER_NUMB...			NUMBER		0	0		
12	QUARTER_OF_YEAR			NUMBER		0	0		
13	CAL_YEAR_NUMBER			NUMBER		0	0		

100%



Defining Cubes

- **Cube is a collection of Measures (Data)**
- **All measures in a cube have the same dimensionality**
- **Use OWB Cube Wizard to build Cubes**
- **Cube can be ROLAP or MOLAP**

Cube: Dimension Order



**Think about sparsity and use of compression first.
(Compression means the use of compressed composites)**

Create Cube

General Implementation Details Rules Summarize To Cache

These settings affect the performance of an analytic workspace in both querying and maintenance processes, such as data loading and aggregation

Dimension Order and Sparsity:

Order	Dimension	Sparse
1	TIME	<input type="checkbox"/>
2	CUSTOMER	<input checked="" type="checkbox"/>
3	PRODUCT	<input checked="" type="checkbox"/>

☒ Use Compression (recommended only for extremely sparse Cubes)

Data Type of Cube: DECIMAL

☒ Partition Cube

Choose a level within a hierarchy of one dimension. One partition will be created for each member of the selected level

Dimension: TIME

Hierarchy: CALENDAR

Level: YEAR

Help Create Cancel

Cube Dimension (Advanced): Compression



- **What is a compressed composite?**
- **When can compression be used?**
- **How sparse is "extremely sparse"?**
- **Rules of thumb**




What Is a Compressed Composite

- **Normal composite has tuples for**
 - ☐ all the leaf values, and
 - ☐ all the precomputed aggregate values (aggindex no), or
 - ☐ all the aggregate values (aggindex yes)
- **With sparse data many aggregate tuples may have only a single child and hence have the same data value as their child.**

Single Child Situation Is Common



- Especially in a multidimensional situation.
- The red nodes can be compressed out.





Compressed Composite (CC) Knows

- **CC knows where these runs of single child parent tuples are.**
- **Stores the common value for these runs only once.**
- **Doesn't materialize the tuples in these runs.**
- **This is fabulous.**
- **Less footprint on disk and in memory, often much less.**
- **Faster aggregation, often much faster.**



CC Limitations in Current Release (10gR1)

- **The only thing you really need worry about is:**
 - ☐ **SUM method or NOAGG method of aggregation only.**

- **Less importantly but in the spirit of full disclosure:**
 - ☐ **No partial aggregation – CC's are so good this doesn't matter (usually).**
 - ☐ **A CC can dimension only a single variable – not a concern to you.**
 - ☐ **A CC's aggregate tuples cannot be updated once built**
 - **To make changes, the aggregates are thrown away.**
 - **CC's are so good this doesn't matter (usually).**



When Can Compression Be Used?

- **SUM method of aggregation**
- **Data are sparse.**
- **How sparse is sparse?**
- **Not as sparse as you might think.**



How Sparse Is Sparse? Use Case #1

- Existing OSA application
- 14 measures
- Time at week, month, year (260 values)
- Product (4,220), customer (7,804) and channel (22)
- Deepish hierarchies on product and customer
- 2.9M input rows
- 9i OSA build on 6Gb. Machine
 - ❑ 616 minutes
 - ❑ 100Gb. on disk



Use Case #1 With AWM10g

- **Slower single cpu machine with 2Gb. Memory**
- **All dimensions in a CC**
- **Partitioned on time at year level**
- **Built in 51 minutes, 1.6Gb. on disk**
- **12x faster, 1/60th of disk**

How Sparse Is Sparse? Use Case #2



TIME at month and year, 10 years, 130 values

CUST: 496,623 values, 2 hierarchies

One is level based with 4 levels

The other is parent-child with depth of 8

SEG: 2 levels, 5 values

RISK: 1,239 values

PRG: 2,658 values

DATA: 11 measures, 31 million input rows

DIMENSIONALITY: Time dense

Should SEG (low cardinality) be in CC or not?

40% dense at least (child and top)



Use Case #2 With AWM10g

- **In 9i:**
 - ☐ Year level data only with skiplevel aggregation.
 - ☐ Took >1 day to load and aggregate.
- **In 10g with AWM10g:**
 - ☐ 1 cpu, 2 Gb. RAM machine
 - ☐ Time dense, other dims in CC.
 - ☐ Partition on time at year level.
 - ☐ No parallelization
 - ☐ 89 min. load & upd. + 115 min. agg = 204 minutes
- **Note: daily load of data would take about 12 or 13 minutes.**
- **With SEG dimension out of the CC aggregation was significantly slower.**

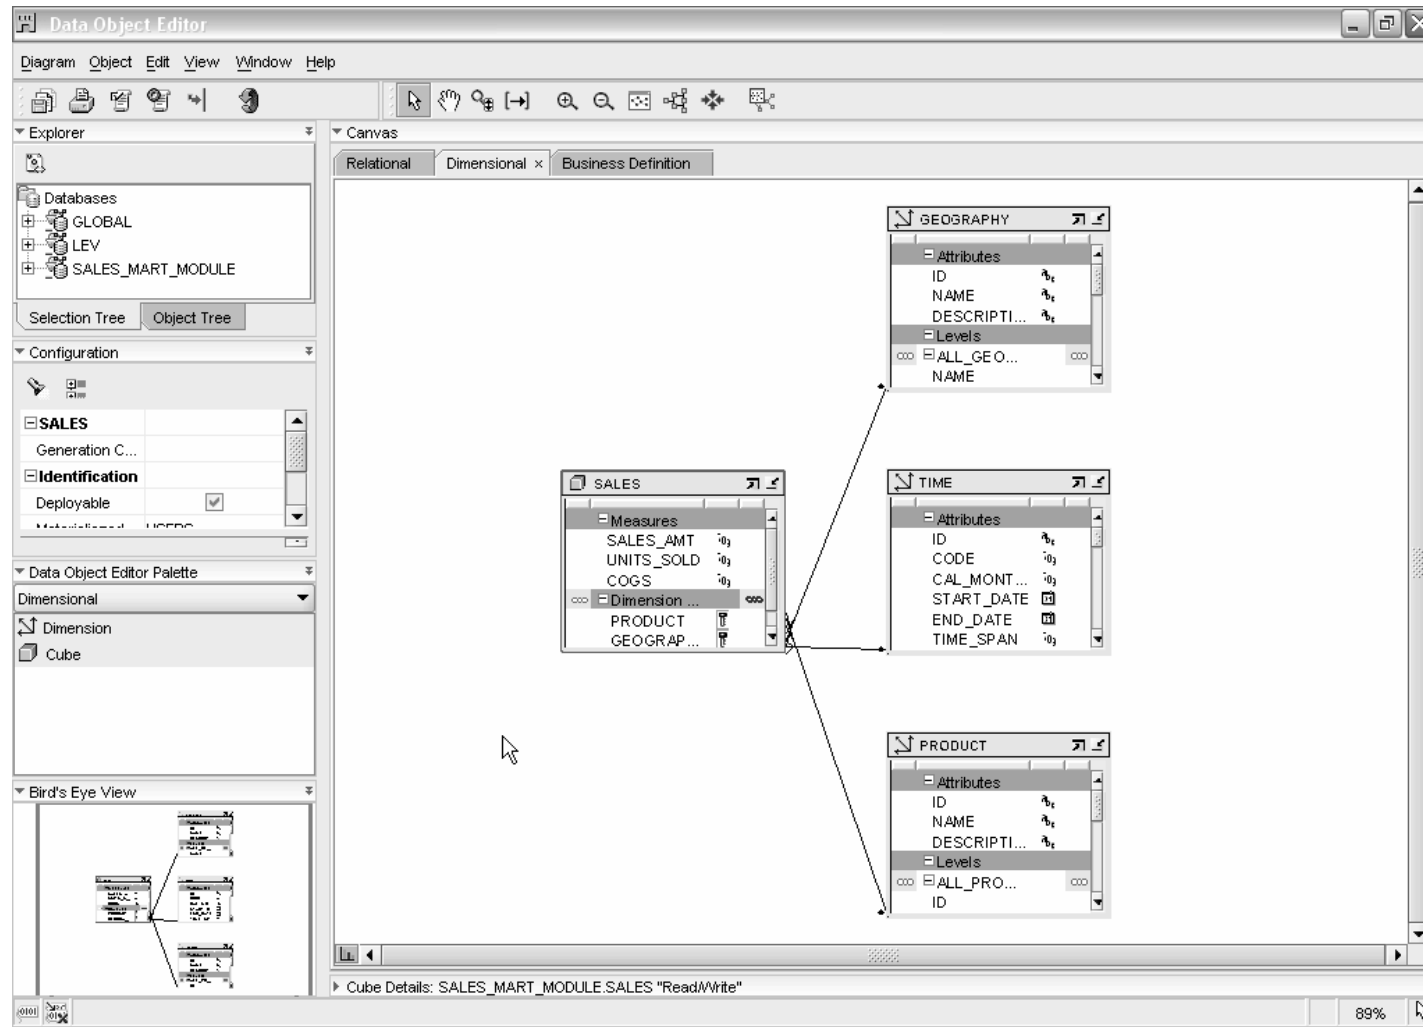


Moral of the story:

- **Our intuition needs to be adjusted.**
- **Experiment with low cardinality dimensions in and out of the CC.**

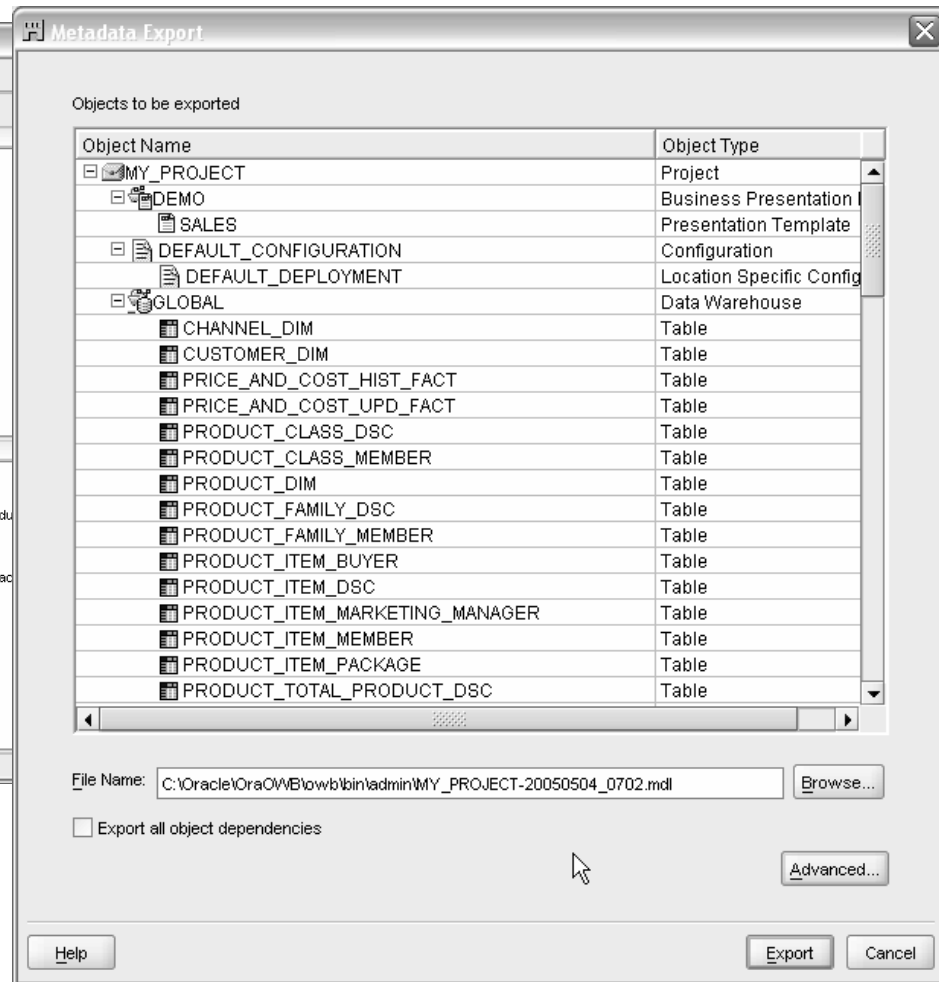
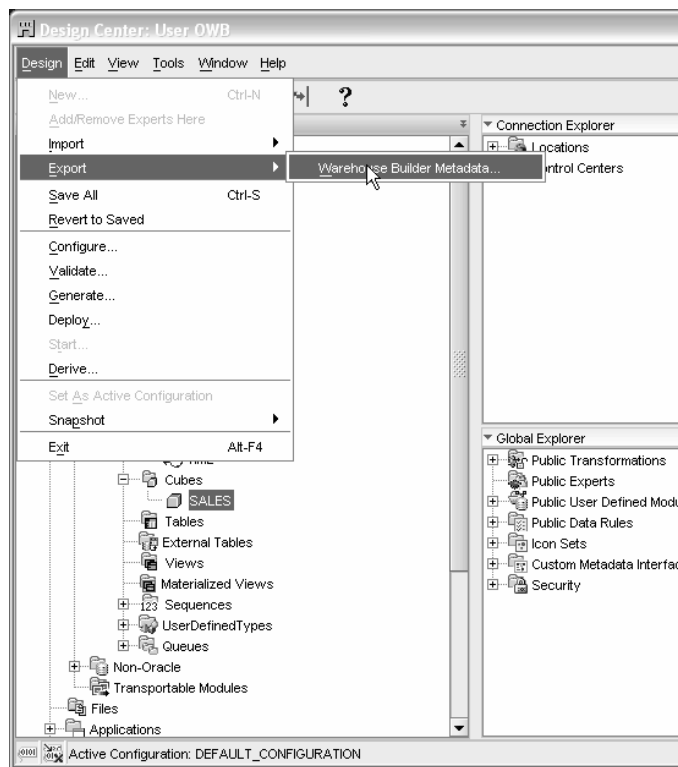


Finished Cube





Deploy







Loading AW Data

- OWB provides Transformations to Load Data into AWs
- Default behavior is to load the entire set of data
- Supports Sub-Setting with customize SQL i.e.

```
procedure ORDERS_LOAD_FILTER
BEGIN
  dbms_awm.create_awcubeload_spec ('ORDERS_FIL', USER, 'ORDERS',
    'LOAD_DATA');
  --- Define the Limiting Where Clause Here
  dbms_awm.Add_AWCubeLoad_Spec_Filter('ORDERS_FIL',USER,'ORD
    ERS',USER,'ORDERS',' month_id>33');
  dbms_awm.refresh_awcube (USER, 'AWS', 'AWORDERS', 'ORDERS_FIL');
  EXCEPTION
    WHEN OTHERS THEN
    NULL;
```

OWB Paris in Action





Managing an OLAP Project

- **Involve end-users early on**
- **Prototype, pilot, then phase 1**
- **Recruit "champion" users**
- **Lead from user community, not IT**
- **Develop in phases**
- **Provide value early on**
- **Keep it simple (at first)**
- **Need forum for users to share ideas**
- **Provide user guide with user's data**



OLAP Implementation Suggestions

- **Pick single first department**
- **Decide on set of terminology at beginning**
- **Use embedded-total objects**
- **Show instances in addition to "levels" in diagrams**
- **Prototype and design iteratively**
- **Pick small initial project. Deliver value quickly**
- **Involve users early on. Listen to feedback**



Conclusions

- **We can finally design OLAP Solutions**
- **Support for both ROLAP and MOLAP (AW)**
- **Strong Foundation for the Future**
- **Still Lacking all the Pieces**
 - ☐ **AWM does not show Mappings (yet?)**
 - ☐ **Manual manipulations in ROLAP or MOLAP cubes not always reflected in OWB metadata**
- **Not Available Yet!**
- **Use OWB 10g R1 NOW in combination with AWM. Works GREAT!**
- **Don't Wait for PARIS!**



Conclusions


- **If you tried or looked at OWB before and said NO. Take another LOOK!**
- **We finally have a full Featured Tool for OLAP end to end design and build!**
- **Lots of new Enterprise Features**
- **Very Low COST!**



How to Get Started?

- **Download OWB 10.1.0.2**
- **Download and install Samples**
- **Read Reviewer's Guide if necessary**
- **Resources:**
 - ☐ **OTN**
 - ☐ **Discussion Forums**


QUESTIONS?



Oracle Warehouse Builder 10g and OLAP – What's New

Oracle OPENWORLD '05

Session #675



Chris Claterbos

claterbos@vlamis.com

Vlamis Software Solutions, Inc.

<http://www.vlamis.com>

Copyright © 2005, Vlamis Software Solutions, Inc.