



Making Your Data Warehouse FASTER

BIWA Summit January 2014

Jonathan Clark
Vlamis Software Solutions
816-781-2880
jclark@vlamis.com
<http://www.vlamis.com>



Vlami Software Solutions, Inc.

- Founded in 1992 in Kansas City, Missouri
- Oracle Partner and reseller since 1995
- Developed more than 200 Oracle BI systems
- Specializes in ORACLE-based:
 - Data Warehousing
 - Business Intelligence
 - Data Transformation (ETL)
 - Web development and portals
- Delivers
 - Design and integrated BI and DW solutions
 - Training and mentoring
- Exclusive supplier world-wide for Windows-based
- Oracle BIC2G BI & EPM VMs
- Expert presenter at major Oracle conferences
- www.vlami.com (blog, papers, newsletters, services)



About Me

- I have worked as both the Office Manager and as a Consultant with Vlami for nearly fourteen years.
- I am responsible for all Amazon Web Services management and training within our team.
- I have worked with clients with a variety of problems using Oracle BI, Oracle OLAP, Application Express, and BI Publisher.
- I love giving training and answering questions.
- I am a regular performer with the Kansas City Renaissance Festival and various historical theatrical events and vaudeville.
- It is real. I use Clubman Mustache Wax. No, I do not use curlers.





Datawarehousing Challenges

- Businesses are integrating business intelligence into every level of their operations
- Datawarehouses are storehouses of historical data organized in a way to provide for the reporting needs of the business
- As the integration of business intelligence grows, so do the demands on the datawarehouse
- How does IT typically respond? Summary tables, materialized views, enormous effort to improve SQL queries and responsiveness, and bigger and better hardware
- Why not OLAP?
 - It is too hard
 - It is incompatible
 - What are you, some kind of weirdo? Nobody uses that.



Definition of OLAP

- OLAP stands for On Line Analytical Processing. That has two immediate consequences: the on line part requires the answers of queries to be fast, the analytical part is a hint that the queries itself are complex.
- i.e. Complex Questions with FAST ANSWERS!



Why use OLAP?

- Empowers end-users to do own analysis
- Frees up IS backlog of report requests
- Ease of use
- Drill-down
- No knowledge of SQL or tables required
- Exception Analysis
- Variance Analysis
- EASY to IMPLEMENT and SUPPORT!



What Does Oracle OLAP Add to a DW?

- 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
- OLAP DML with what-if, forecasting
- Platform for extensions



OLAP Option – High-level View

- Advanced analytics
- Integrated in RDBMS
- Easy to develop
- Easy to use
- Facilitate collaboration
- Flexible deployment
- Scaleable and performant
- True Relational – Multidimensional database




Why OLAP for BI?

- BI often presents data dimensionally
- Dimensions are natural way to look at data
 - By, across, over, time, geography, product
 - Comparison of multiple dimension values
- Multi-dimensional storage of data speeds analysis
- Natural to express dimensional comparisons
 - Share of parent
 - Compared to last year
- Allows for hierarchical dimensions with multiple levels
 - E.g. by country, drill to state, drill to city



Oracle OLAP

Leveraging Core Database Infrastructure




- Single RDBMS-MDBMS process
- Single data storage
- Single security model
- Single administration facility
- Grid-enabled
- Accessible by any SQL-based tool
- Embedded BI metadata
- Connects to all related Oracle data



Oracle OLAP

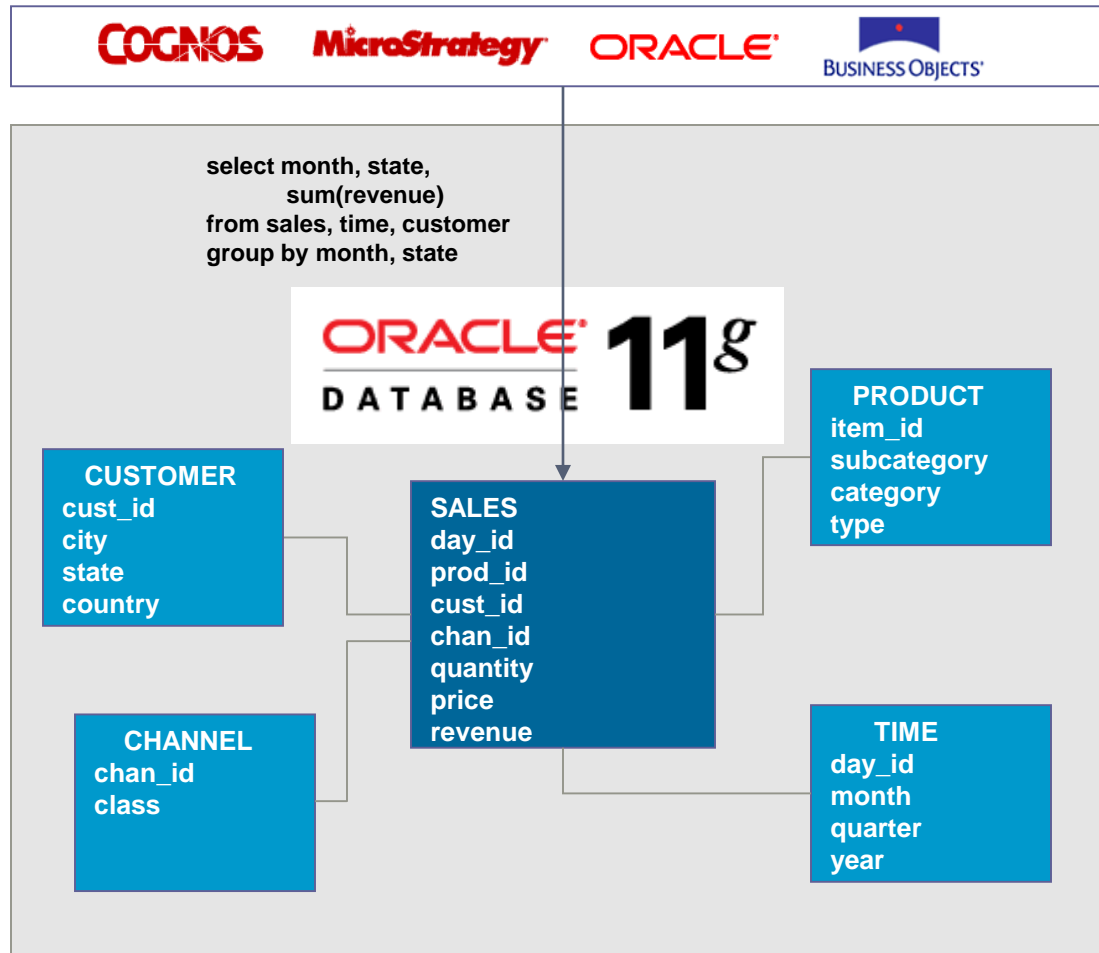
- A summary management solution for SQL based business intelligence applications
 - An alternative to table-based materialized views, offering improved query performance and fast, incremental update
- A full featured multidimensional OLAP server
 - Excellent query performance for ad-hoc / unpredictable query
 - Enhances the analytic content of Business intelligence application
 - Fast, incremental updates of data sets





Materialized Views

Typical MV Architecture Today

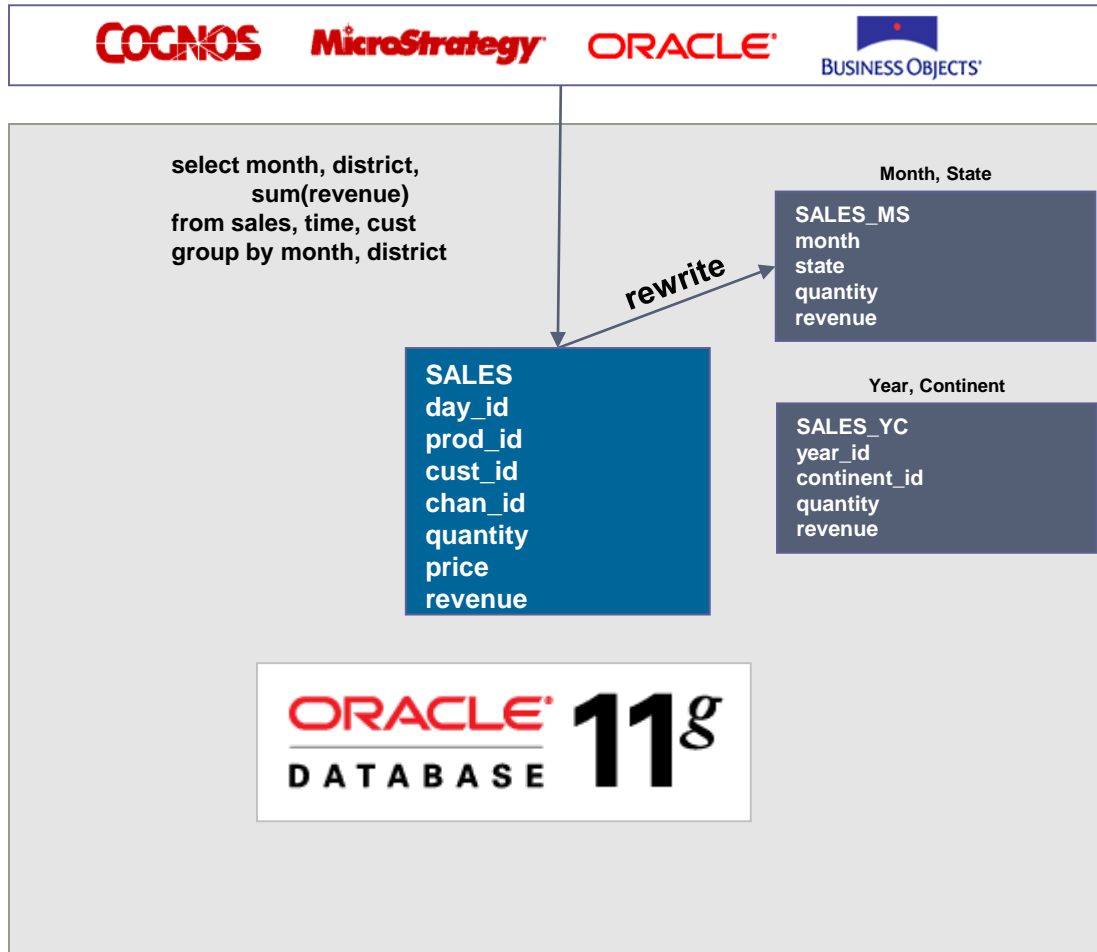


- Query tools access star schema stored in Oracle data warehouse
- Most queries at a summary level
- Summary queries against star schemas can be expensive to process



Materialized Views

Automatic Query Rewrite



- Most DW/BI customers use Materialized Views (MV) today to improve summary query performance
- Define appropriate summaries based on query patterns
- Each summary is typically defined at a particular grain
 - Month, State
 - Qtr, State, Item
 - Month, Continent, Class
 - etc.
- The SQL Optimizer automatically rewrites queries to access MV's whenever possible



Materialized Views

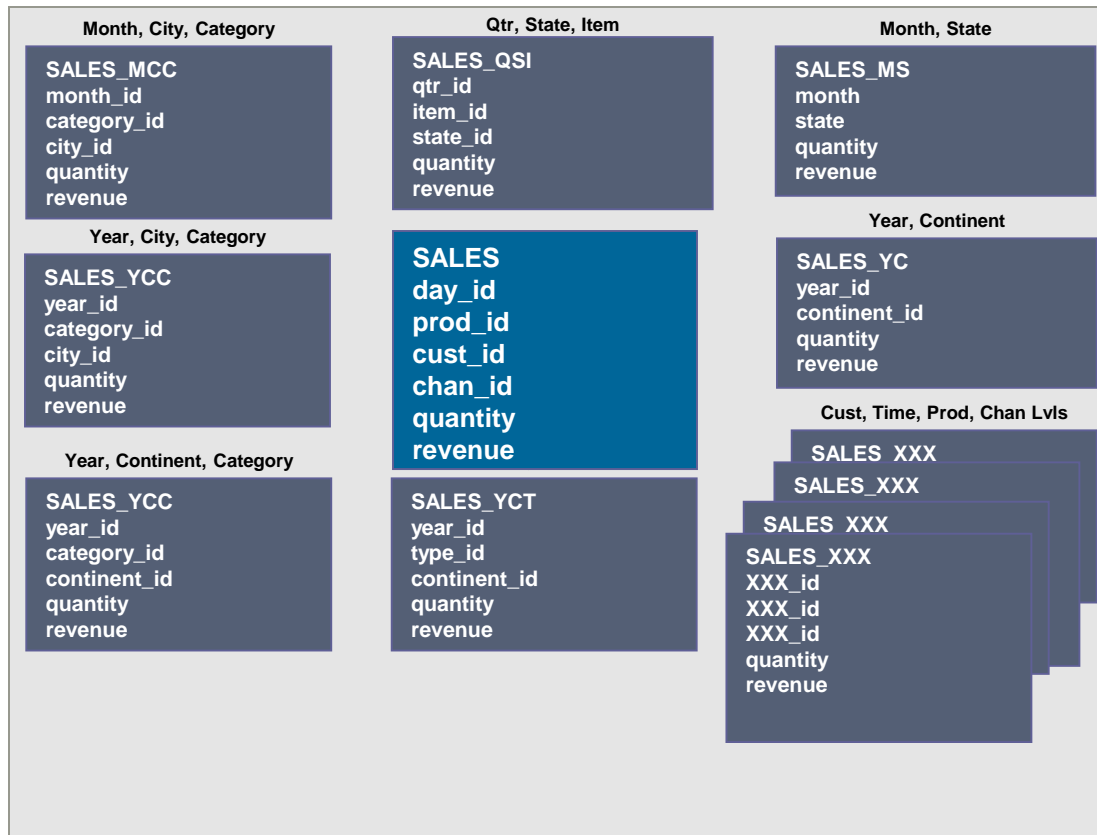
Challenges in Ad Hoc Query Environments

Cognos

MicroStrategy

ORACLE

BUSINESS OBJECTS




- Creating MVs to support ad hoc query patterns is challenging
- Users expect excellent query response time across any summary
- Potentially many MVs to manage
- Practical limitations on size and manageability constrain the number of materialized views



Cube-based Materialized Views

Much Better Manageability & Performance




- A single cube provides the equivalent of *of thousands of summary combinations*
- The 11g SQL Query Optimizer treats OLAP cubes as MV's and rewrites queries to access cubes *transparently*
- Cube refreshed using standard MV procedures



Cost Based Aggregation

Pinpoint Summary Management



- Improves aggregation speed and storage consumption by pre-computing cells that are most expensive to calculate
- Easy to administer
- Simplifies SQL queries by presenting data as fully calculated



Easy Analytics

Fast Access to Information Rich Results

- Time-series calculations
- Calculated Members
- Financial Models
- Forecasting
 - Basic
 - Expert system
- Allocations
- Regressions
- Custom functions
- ...and many more

Snapshot of some functions

deprdecl	aggregate	abs	rank	chgdims
deprdeclsw	allocate	antilog	rem	instat
deprsl	categorize	antilog10	remainder	limit function
deprsoyd	correlation	arccos	round	statall
fintsched	fcopen	arcsin	sign	statdepth
fpmtsched	fcquery	arctan	sin	statequal
growrate	info	arctan2	sinh	statfirst
irr	normal	bin_to_num	smooth	statlast
npv	random	bitand	sort	statlen
vintsched	stddev	ceil	sqrt	statlist
vpmtsched	any	cos	tan	statmax
cumsum	average	cosh	tanh	statmin
lag	count	decode	truncate	statrank
lagabspect	every	exp	width_bucket	statval
lagdif	largest	floor	begindate	coalesce
lagpct	median	greatest	dayof	na2
lead	mode	intpart	ddof	nafill
movingaverage	none	least	enddate	naflag
movingmax	percentage	log function	endof	nullif
movingmin	smallest	log10	isdate	nvl
movingtotal	forecast	max	makedate	nvl2
total	modulo	min	mmof	ascii




Easy Analytics

Optimized Data Access Method

How do Expenses compare this Quarter versus Last Quarter

What is an Item's Expense contribution to its Category?



- Data stored in dense arrays
- Offset addressing – no joins
- More powerful analysis
- Better performance



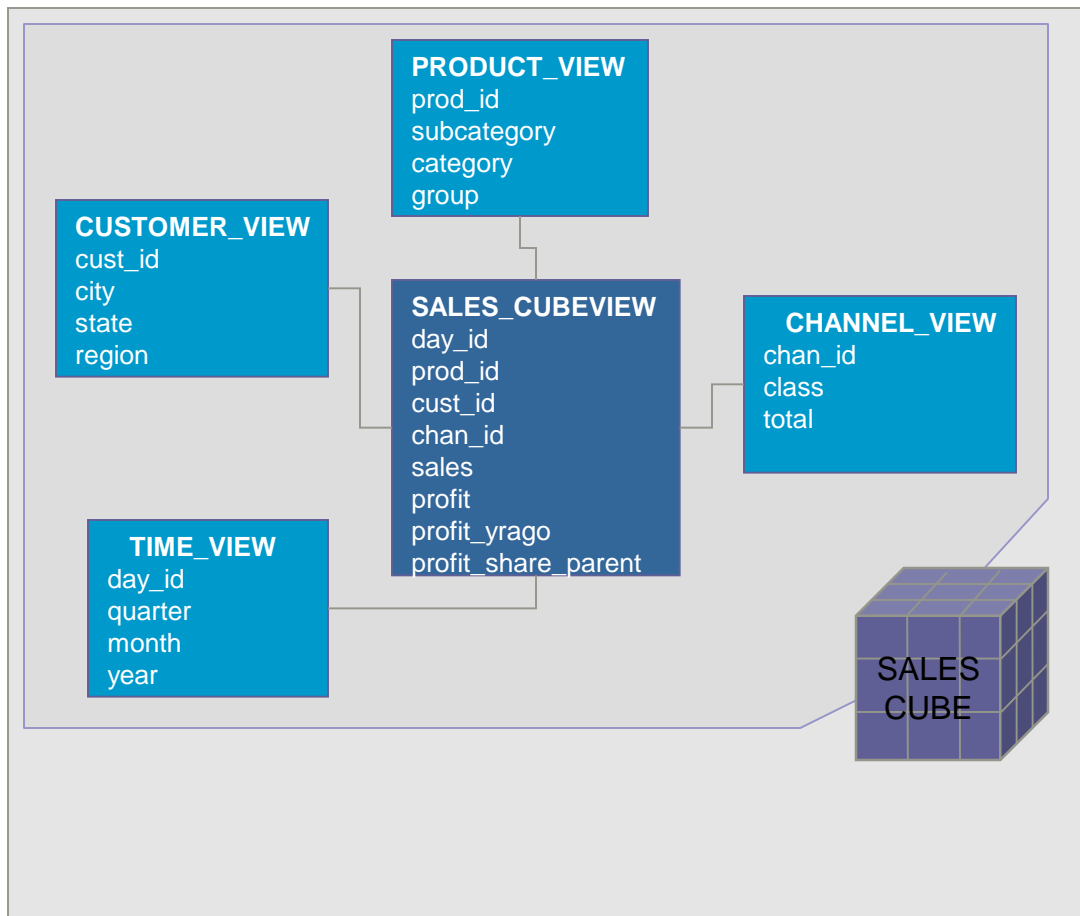
One Cube Accessed Many Ways...

- One cube can be used as
 - A summary management solution to SQL-based business intelligence applications as cube-organized materialized views
 - A analytically rich data source to SQL-based business intelligence applications as SQL cube-views
 - A full-featured multidimensional cube, servicing dimensionally oriented business intelligence applications



Cube Represented as Star Model

Simplifies Access to Analytic Calculations



- Cube represented as a star schema
- Single cube view presents data as completely calculated
 - Analytic calculations presented as columns
 - Includes all summaries
- Automatically managed by OLAP




Empowering Any SQL-Based Tool

Leveraging the OLAP Calculation Engine

Application Express on Oracle OLAP

```
SELECT cu.long_description customer,  
       f.profit_rank_cust_sh_parent,  
       f.profit_share_cust_sh_parent,  
       f.profit_rank_cust_sh_level,  
       f.profit,  
       f.gross_margin  
FROM time_calendar_view t,  
     product_primary_view p,  
     customer_shipments_view cu,  
     channel_primary_view ch,  
     units_cube_view f  
WHERE t.level_name = 'CALENDAR_YEAR'  
      AND t.calendar_year = 'CY2006'  
      AND p.dim_key = 'TOTAL'  
      AND cu.parent = 'TOTAL'  
      AND ch.dim_key = 'TOTAL'  
      AND t.dim_key = f.TIME  
      AND p.dim_key = f.product  
      AND cu.dim_key = f.customer  
      AND ch.dim_key = f.channel;
```






Oracle OLAP Summary

- Improve the delivery of information rich queries by SQL-based business intelligence tools and applications
 - Fast query performance
 - Simplified access to analytic calculations
 - Fast incremental update
 - Centrally managed by the Oracle Database




Building Cubes in AWM






Viewing Cubes in Excel





QUESTIONS?





Making Your Data Warehouse FASTER

BIWA Summit January 2014

Jonathan Clark
Vlamis Software Solutions
816-781-2880
jclark@vlamis.com
<http://www.vlamis.com>