

IS OLAP DEAD? Can Next Gen Tools Take Over?

1 - - -

×

GoToWebinar Control Panel





Presentation Slide Deck http://www.senturus.com/resources/



RESOURCE LIBRARY

An extensive, free library of past webinars, demonstrations, whitepapers, presentations, helpful hints, and more.



Agenda

- Introductions
- OLAP Defined
- Why OLAP Pros/Cons
- Current State of OLAP Architectures
- New Visualization Tools
- Fast Columnar/In-Memory Databases
- Big OLAP on Big Data
- Senturus Overview
- Additional Resources
- Q&A



Introducing...Today's Presenters





Pedro Ining Senior BI Architect Senturus, Inc.

Jim Frazier VP of Client Services Senturus, Inc.



Poll

What is the state of OLAP in your organization today?

- OLAP is going strong
- We have already begun moving away from OLAP
- We are committed to both OLAP and the new technologies





OLAP DEFINED

Is OLAP Dead?

Quotes on the state of OLAP Cubes

- "In-memory databases are killing OLAP."
- "The OLAP cube is history."
- "Is OLAP still relevant?"
- "Building cubes for Tableau can be a waste of time!"



OLAP, OLTP, DWH Defined

OLTP - Online Transaction Processing

- Core data repository of data flowing into business systems (ERP)
- Optimized for quick data entry and relational integrity
- Not optimized for reporting and data analysis
- Typically very complex schema design with many normalized tables that facilitate high volume throughput of transactions

DWH - Data Warehouse

- Central repository of multiple ERP systems or other data sources
- The data warehouse (DWH) is a database system separate from the OLTP
- Architected in an optimized fashion for easy reporting and analysis
- Organizes and hides the complexity of the OLTP systems for efficient, timely and accurate reporting



OLAP, OLTP, DWH Defined

OLAP - Online Analytical Processing

- Edgar F. Codd 'father of relational database' coined the term OLAP, Arbor/Essbase went on to market the term
- The data warehouse plus a **central repository** that defines the relationships between tables (facts/dimensions) and stores complex business rules/calculations (e.g. YTD, YTD LY, Margins, Inventory Turns, etc.)
- Allows for **high performing** interactive analysis
- Generally referred to as 'cubes'



Modern BI Architecture



Typical Best Practices BI System with OLAP Layer Architected Data Warehouse & BI Tools Microsoft +ableau **Properly Staged Data BI Tools** Single Version of the Truth Abstraction Standard Threshold Reports or Desktop Viz Tools Alerting Model Data Marts/Warehouse Ad Hoc Querying Data Self-Service Reporting & Slicing & Analysis **OLAP** Dicing Layer Web Portal Dashboard Authoring Conforming Dashboards/ **Business Process Scorecards Dimensional Models*** Report Authoring Information Security Thresholdbased Alerts



* Also known as a Star Schema

Why OLAP

- Historically, size and speed limitations of databases limited query performance
- Central repository for relationships and complex business calculations
- Buffers the business user from complex native database structures and sensitive calculation logic
- Cubes generally have higher performance vs. relational queries
- Fast, simple, drag-and-drop ad-hoc analysis and reporting
 - Self-Service with guardrails
- Visual exploration
 - Multi-dimensional view of data
 - Drill-down on hierarchies
- Many business users love the interface and are used to querying by governed data dimensions and measures that are prebuilt for them



OLAP Familiar Interfaces





Why Not OLAP

- Massive increase in data volumes
 - Latency large cubes increase cube build times, impacting SLAs
 - Large cardinality dimensions and many dimensions
 - Real-time updates are difficult if not impossible
- Movement of data into another proprietary structure
- Upfront investment in cube modeling
 - Measures, dimensions, hierarchies all defined upfront
 - Not a flexible agile BI environment
 - New cube builds and designs are required as the business changes
- Continued developer maintenance and administration
- CPU power, memory and powerful servers are very affordable do we still need the OLAP layer





The Current State of OLAP Architecture Traditional OLAP

TRADITIONAL OLAP ARCHITECTURES

MOLAP - Multi Dimension OLAP

- Most traditional OLAP design
- Data is stored in the multidimensional cube
- Data is moved from the relational database to the cube
- Data is pre-aggregated and allows for very fast analysis

ROLAP - Relational OLAP

- Modeled on top of the relational star schema database
- Data storage is kept in the relational database
- Utilizes SQL to query the DB in an OLAP manner
- May use proprietary in-memory caching techniques

HOLAP - Hybrid OLAP

- Combines the advantages of MOLAP and ROLAP
- Stores summary data in MOLAP structure
- Can "drill-through" to relational database for more detail



Top OLAP Products

For Dimensional BI Uses

- IBM Cognos Transformer Cubes (MOLAP)
- Microsoft SQL Server Analysis Services (SSAS)
 - Dimensional and Tabular (MOLAP/HOLAP)
- IBM Cognos Dynamic Cubes (ROLAP)
- MicroStrategy (ROLAP)

Typically For Finance Use

- IBM Cognos TM1 (writeback)
- Hyperion Essbase (writeback)



COGNOS POWERPLAY (TRANSFORMER)

Advantages

- Performance (vs. relational)
- Easy to use and develop
- ETL-like capabilities (limited) i.e. no star schema needed
- Can act as meta-data layer
- Great relative-time calc capabilities (YTD, Rolling 13 months)
- Less intensive hardware requirement

Challenges

- Significant cube size limitations
- Limited categories per dimension level
- Cube builds take time and cubes exist as separate files (.mdc)
- Lacks capabilities now available in other OLAP tools
- Row-level (dimensional) security is very challenging to maintain
- Unclear product support going forward
- Only works in the IBM Cognos stack



IBM COGNOS DYNAMIC CUBES



- IBM Cognos Dynamic Cubes was added to the Cognos 10.2 BI suite as an inmemory relational OLAP product that could address the challenge of highperformance/low latency interactive analysis against terabytes of data
- The last significant update to Dynamic Cubes occurred in version 10.2.2; IBM has since focused most development efforts on the Cognos Analytics v11 release
- Currently, no current plans by IBM to enhance the Dynamic Cubes product



IBM COGNOS DYNAMIC CUBES

Advantages

- Scalability limited only by database and RAM cache sizing
- Handles large dimensions well, allows dimension attributes
- Built-in relative time calcs on par with Transformer
- MDX Scripting can set up just about any type of calculation
- Dynamic security can set up dimensional filtering so that all security is derived from SQL tables
- Aggregate aware can dynamically select database aggregate tables or inmemory aggregates for fast results

Challenges

- Requires star or snowflake schema as data source
- Cache needs to be warmed for decent performance
- Requires 64-bit application server and may require significant memory footprint for large cubes (e.g. 64-128GB)
- Report authors require dimensional reporting experience
- CAN ONLY BE USED BY COGNOS BI STACK (Senturus has developed the Analytics Connector to access Tableau)



Dynamic Cubes in Play

Large Health Insurance Provider Deployed Dynamic Cubes





- Finance project used IBM Cognos Dynamic Cubes to replace legacy Cognos Transformer cubes, went into production Q1 2017
- Large number of reports were converted or created on top of the Dynamic Cube to provide a guided set of highly formatted reports that allowed drill-down
- Many complex business calculations were developed and stored in the cube, report writers can leverage a central set of calculations without having to write them in the report

MICROSOFT SQL SERVER ANALYSIS SERVICES (SSAS)





Multi-Dimensional	Tabular
Dimensions and Measure Groups	Tables and Relationships
Highly Scalable and Mature	Fast Design and In-Memory
Feature Rich and Complex	Easy to Get Started



MICROSOFT SQL SERVER ANALYSIS SERVICES(SSAS) TABULAR MODEL

- Introduced in SQL Server 2012
- Model paradigm = tables and relationships
- Data stored in-memory



SQL Server Analysis Services Tabular Model

- Uses a different engine (xVelocity) and uses a columnar DB structure
- Combines the functionality of MOLAP cubes and relational DBs



MICROSOFT SQL SERVER ANALYSIS SERVICES (SSAS) TABULAR

Advantages

• Simpler data development model, faster to develop



SQL Server Analysis Services Tabular Model

- Generally much faster than MOLAP
- DAX learning curve is easier than MDX
- Fast COUNT DISTINCT queries

Challenges

- Dependent on server memory footprint (DirectQuery mode available in 2016)
- Some multidimensional features are not available (e.g. many-to-many)
- Complex calculations may be difficult to implement
- Large datasets



SSAS in Play

Major American Clothing Company

- Re-architected an older Oracle based data warehouse to a SQL Server
- User community already very familiar with cube technologies
- Wanted to use SSAS OLAP cubes for their advanced relative time calcs
- Ability to create complicated advanced inventory calcs and on the fly currency conversions
- Ability to set defaults for certain dimensions such as currency type
- SSAS fits into their corporate strategy for multiple tools
- SSAS Tabular was chosen for performance and flexibility





New Generation Visualization Tools

New Generation Visualization Tools







- Over the last few years desktop visualization tools have sprouted on desktops throughout the enterprise
- IBM Cognos Analytics v11 allows similar functionality over a web interface
- Rich visualizations are easily created by business users without the help of IT
- Decentralized model of data governance
- No waiting on developers to create next iteration of an OLAP cube
- Allows users to integrate data on the desktop/web
- Creation of desktop 'micromodels' (Tableau data extracts)
- Can use OLAP datasources, but works best with non-OLAP sources
- Can begin to have performance issues when creating large data extracts or going against large datasources



Tableau



\$ ← → 🖬 🛱	G	• O •	4.		t s⊅ tB	18 1	· // · LL \$4	Standard 🝷 🔝 🖞 🖓			
Analytics	٥	Pages			iii Columns	iii Columns Goals					
Local_NHL_TC17 (Local					I≣ Rows	🗆 Le	ague	Conference Division Team	1		
Dimensions D	*	Filters			Motadata & Navigation						
Position	^				Metauata & Navigation						
E Round					League	Confere	Division	Team			
5 Shoots		Marke			National	Eastern	Atlantic	Boston Bruins			
State Province		III Bar 🔻			Hockey League			Buffalo Sabres			
Teams								Detroit Red Wings			
A Teams 2016-2017			Ð	T				Florida Panthers			
♣ ¹ League		Color	Size	Label				Montreal Canadiens			
2 Conference								Tampa Pay Lightni			
♣³ Division		Detail	Toolitin					Toronto Maple Leafs			
å4 Team		Detail	Toolup				Matropolitan	Carolina Hufricanes			
Conference							Metropolitali	Columbus Blue Jac			
Division								New Jersey Devils			
E League								New York Islanders			
💷 Team								New York Rangers			
Team ABR								Philadelphia Flyers			
Time								Pittsburgh Penguins			
Season 2016-2017								Washington Capita.			
Date Day						Western	Central	Chicago Blackhawks			
Date Name								Colorado Avalanche			
Date Short								Dallas Stars			
Month Name								Minnesota Wild			
easures								Nashville Predators			
Goals	^							St. Louis Blues			
Assists								winnipeg Jets			

OLAP concepts like dimensions and hierarchies are alive and well in Tableau



Tableau Data Extracts

"Building Cubes for Tableau can be a waste of time!"

- TDE is a compressed snapshot of data stored on disk and loaded into memory as required
- Data engine can be described as its own "in-memory analytic database"
- Stores data in a columnar store structure
- Dramatically reduces the input/output time required to access and aggregate values
- New in upcoming 10.5 hyper data engine

Reasons to Use TDEs

- Better performance vs. connected data sources
- Reduced load on connected data sources
- Portability can be bundled in a packaged workbook for easy sharing
- Pre-aggregation option to aggregate data for visible dimensions "aggregated extract"



Tableau and Cubes

"But wait Tableau does support cubes."

- Leverages
 - One version of the truth
 - Complex calcs that are already created in the cube
 - Billions of rows response time across pre-aggregated calcs faster
- Tableau will work if you stay within the structure of the cube
- Functional differences
 - No cube extracts (10.4 supports BW cubes)
 - No user-defined hierarchies
 - Aggregations controlled in the cube
- Supports
 - Oracle Hyperion Essbase
 - Teradata OLAP
 - Microsoft Analysis Services (SSAS)
 - SAP NetWeaver Business Warehouse
 - Microsoft PowerPivot
 - Analytical Views in SAP Hana





IBM Cognos Analytics v11

- Cognos Analytics v11 architecture adds data modules which represent a major shift in the central metadata layer (framework) paradigm
- Data modules allows end users to quickly add new data sources and model new data subjects without having to wait for DWH changes
- Uploaded files and data sources can be stored as 'snapshots' on the server's file system using the Apache Parquet columnar file storage mechanism
- Allows for fast query response times





FAST COLUMNAR AND IN-MEMORY

Databases

"IN-MEMORY DATABASES ARE KILLING OLAP"





COLUMNAR & IN-MEMORY DATABASES

Columnar Databases

- Traditional databases store data by each row
- Columnar databases store data in columns rather than in rows
- This storage architecture can result in high-performing queries especially aggregation queries
- Example DBS:
 - Sybase IQ
 - IBM DB2 with BLU Acceleration
 - A capability built into DB2, not a separate install component
 - Focus on analytics
 - Dynamic in-memory, does not require all data to be in-memory
 - Columnar and traditional row-based tables
 - SQL Server 2014/16
 - Columnar store indexes
 - In-memory OLTP tables



35

BUT IF REMOVE THE OLAP LAYER

- Raw queries will be fast, but what about the semantic layer?
- You could use relational models with some level of metadata and calculations
- But complex calcs, dimensions, drill downs would be missing





THE CURRENT STATE OF OLAP ARCHITECTURE

The New School Big OLAP on Big Data

Big OLAP on Big Data

Several new vendors and open source solutions are creating new scalable OLAP on Hadoop products











Big OLAP on Big Data

- Slow performing queries on big data implementations are driving new OLAP technologies
- Classic OLAP technologies on big data necessitated movement of Hadoop data into traditional relational data warehouses further increasing latency
- New OLAP technologies are architected to be part of the Hadoop stack and allow queries across Hadoop with no additional movement of data



Big OLAP on Big Data







SUMMARY WRAP-UP

So Is OLAP Dead?

- In general, the key concepts of OLAP dimensions, measures, hierarchies, drill-down are still alive and well but the technology that surfaces those concepts are changing
- Business users will always want a high performing BI layer that is easy to use and allows for interactive BI
- Some will will want a central repository that contains all the relationships, hierarchies and complex business rules already developed
- Other users like data scientists and advanced business analysts will want a more agile free form solution with have high performance



BI Architectural Guidance

- Best practices
- Pragmatic recommendations
- Path to an architecture that supports current needs and grows for future

Consultation with Senturus Senior BI Architect
at special prices for webinar attendees20 hours\$379940 hours\$7399





WHO WE ARE Business Analytics Consultants

Bridging the Gap Between Data & Decision Making





Business Analytics Architects

- Dashboards, Reporting & Visualizations
- Data Preparation & Modern Data Warehousing
- Self-Service Business Analytics
- Big Data & Advanced Analytics
- Planning & Forecasting Systems
- Proprietary Analytics Connector Software







1000+ Clients, 2000+ Projects, 16+ Years





ADDITIONAL RESOURCES

Upcoming Events

http://www.senturus.com/events/



ENHANCING BI WITH PREDICTIVE ANALYTICS

REAL WORLD APPLICATIONS

Thursday, October 26, 2017 - 11am PT/ 2pm ET - 60 minutes



TABLEAU LIKE A PRO IN HALF A DAY

HANDS-ON WORKSHOP IN AUSTIN, TX

Wednesday, November 8, 2017 - 8:30am – 12:30pm CT or 1pm – 4:30pm CT - 4 hours or 3.5 hours



TABLEAU LIKE A PRO IN HALF A DAY

HANDS-ON WORKSHOP IN PALO ALTO, CA

Wednesday, November 15, 2017 - 8:30am – 12:30pm PT or 1pm – 4:30pm PT - 4 hours or 3.5 hours



Other Resources





RESOURCE LIBRARY

BLOG

http://www.senturus.com/senturus-resources/



Cognos & Tableau Training Options



TABLEAU	10	DESKTOP	INTERMEDIATE DATA VISUALIZATION AND DASHBOARDING	ост 13	•
TABLEAU	10	DESKTOP	ADVANCED DATA VISUALIZATION AND DASHBOARDING	ост - ост 16 17	•
IBM COGNOS	11	COGNOS ANALYTICS	OVERVIEW, NEW FEATURES AND BUSINESS AUTHORING	ост 16	•
IBM COGNOS	11.0.5	COGNOS ANALYTICS	SELF-SERVICE DATA MODULES AND DASHBOARDING IN V11	ост 17	•
IBM COGNOS	11.0.5	SYSTEM ADMINISTRATION	SYSTEM ADMINISTRATION	ост - ост 18 19	•

NTURUS



Q&A

If your question or issue is broader than what can be answered today, contact us at

info@senturus.com

and we will set up a free consultation.





Thank You!

www.senturus.com info@senturus.com 888 601 6010



Copyright 2017 by Senturus, Inc. This entire presentation is copyrighted and may not be reused or distributed without the written consent of Senturus, Inc.



