

Today's Data Warehousing Challenge

Based on a set of relatively mature practices and technologies, relational data warehousing is a foundational component of enterprise data management strategies. However, many mid-market businesses and divisions of large enterprises have never undertaken effective data warehousing initiatives, due to challenges with data integration and governance on the one hand, and identifying the right set of reports and dashboards to produce as the core set of information delivery interfaces on the other. The proliferation of data sources beyond the traditional ERP and line of business systems into a myriad of marketing, sales, and service processes and communications channels on the customer side, and the ability to collect vast amounts of data from systems that produce and deliver goods and services (together often referred to as “big data”) has somewhat stymied the adoption of data warehouses by increasing the scope and complexity of the integration challenge.

Introducing the EC Wise Solution: Decision Hub

The good news is that tools are evolving along with the need. Into this environment, EC Wise introduces its DecisionHub service, a framework for rapid implementation of industry specific data warehouses and marts, with out of the box integration with data feeds from social networks , one to one marketing, loyalty programs, and offer generation/management systems. We build DecisionHubs on the market leading Oracle 12c database, using Talend integration components, and Jasper dashboards and reports, with database schema designed by industry domain experts and pre-built industry specific dashboards and reports. The first industry specific version of DecisionHub was developed and is being marketed by EC Wise partner Engage Innovations for the casino gaming industry. We expect to bring versions for general hospitality, spend management and security to market in the very near future.

We can deliver DecisionHubs as pure software or preconfigured on the [Oracle Database Appliance](#) (ODA). The database engine, Oracle 12c is installed on the base system, and two virtual machines are included, one running the Talend integration services, the other running the Jasper reports and dashboard. Engineering and pre-tuning the system for a particular set of workloads eliminates the necessity of integrating services and tuning resource allocation between them for each customer deployment. This saves a significant amount of deployment cost, and significantly accelerates time to value.

The Database: Oracle 12C

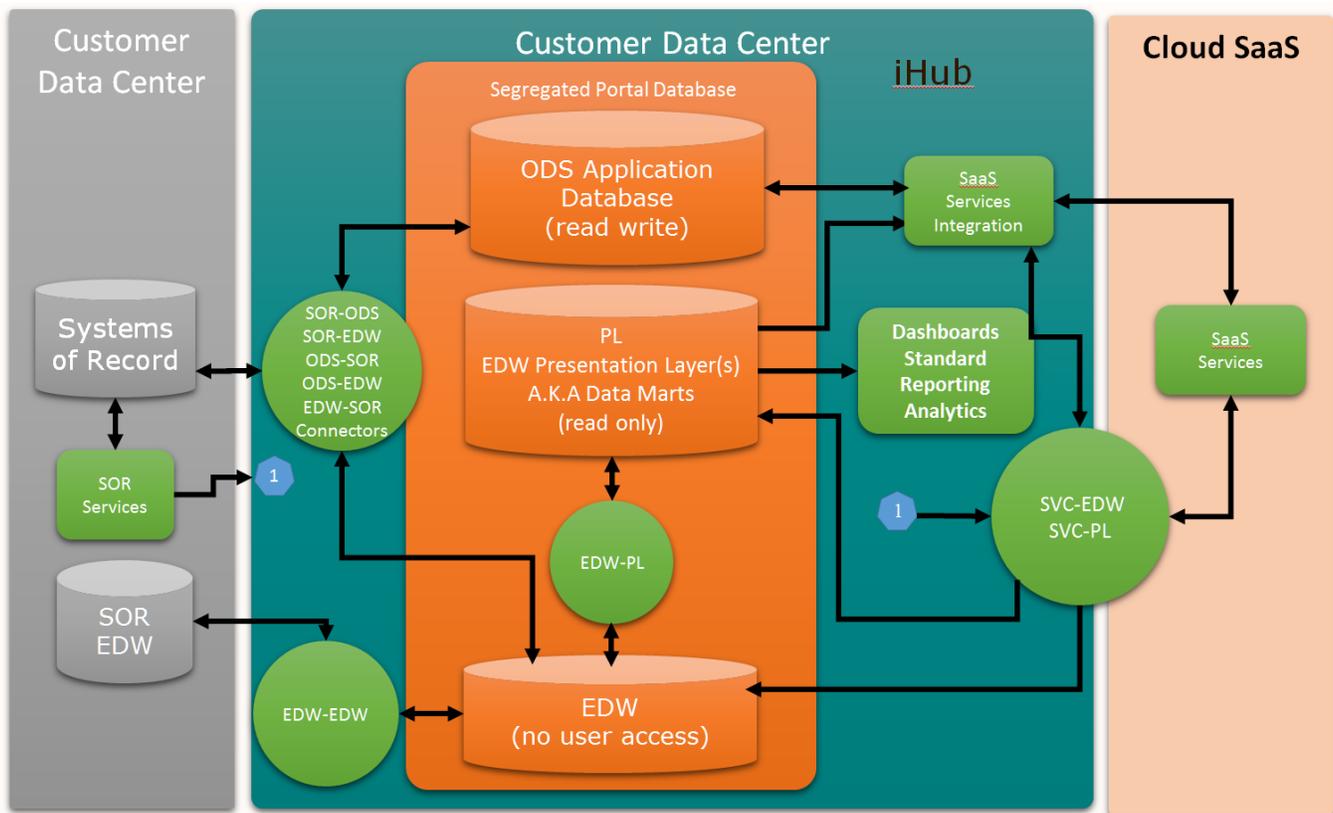
Oracle elaborates on some of the issues involved in data warehousing today in their 2014 white paper, “[Oracle Database 12c for Data Warehousing and Big Data](#)”. That white paper also highlights the Oracle Database 12c capabilities that are true enabling technologies for large scale data warehousing. These include:

- Oracle’s “star transformation” leverages bitmap indexes on the fact table to efficiently join multiple dimension tables in a single processing step. EC Wise has designed the DecisionHub schema to take full advantage of this powerful technology.
- Partitioning is the feature that allows a database to scale for very large datasets while maintaining consistent performance, without needing to increase hardware resources. Oracle leads the industry with

the most comprehensive set of partitioning technologies, offering numerous methods for partitioning tables, including support for custom partitioning schemes. With partitioning enabled, we can partition DecisionHub fact tables on one or two key dimensions, such as customer and activity date.

- Oracle parallel execution enables the database engine to take advantage of the ample CPU and IO resources in the ODA, particularly when applied to partitioned tables and indexes.
- Oracle's partition pruning and partition-wise joins take advantage of partitioned tables, and involve making a first query pass that identifies the partitions that contain qualifying data, and only including those partitions in the query plan.
- In 12c, Oracle introduced Adaptive Query Optimization, which enables the engine to make run-time adjustments to execution plans, based on any changes in the data landscape.
- Oracle's sophisticated resource manager allows administrators to set up rules about how much of the system's resources each group of resource consumers can use. This allows DecisionHub to support multiple types of workloads from other systems querying the warehouse to prepare marketing campaigns, to analysts doing ad hoc reporting and managers viewing up to the minute business results on their dashboards.
- Oracle Database In-Memory is a transparent enhancement, which when enabled causes selected columns to be loaded into an in-memory Column Store which Oracle scans with Single Instruction processing Multiple Data values (SIMD) vector processing. With Database In-Memory, Oracle has introduced a new aggregation algorithm for the join-and-aggregate operations found in typical data warehouse queries.

Taken together, the database engine enhancements Oracle has introduced in versions 11 and 12 make it possible to run multiple classes of workload on a single database instance, and allow DecisionHub to deliver high performance results on a database where loading, reporting, application queries and analytics all take place simultaneously, *without needing to have a special OLAP engine installed!*



DecisionHub Reference Architecture

DecisionHub Integration: Built on Talend

Using Talend’s integration platform in conjunction with Oracle Database 12c makes DecisionHub an incredibly flexible and high performance platform. Talend was designed for a variety of use cases including ETL to data warehouse processes, and provides capabilities that support DecisionHub’s metadata and governance features. With each installation of DecisionHub, customers get the Talend Administration Center, which provides system monitoring, scheduling and control over integration processes.

The Platform: Oracle Database Appliance

Oracle’s ODA is a fully integrated and redundant system of software, servers, storage and networking that delivers high-availability database and virtualization services for a wide range of OLTP and data warehousing applications. By delivering DecisionHub on ODA we minimize the effort required to deploy and configure system services on a platform optimized for a combination of database and virtualization workloads. With ODA, we can provide end customers a unique database license scaling model, whereby they can license the database to run on a subset of the machine’s 36 processors (in the X5-2) and add processors to their license as their needs grow.

Marketplace Results

After EC Wise partner Engage Innovations delivered its first DecisionHub, EC Wise CEO Jack Hakim commented that “While it took a great effort from our development and IT teams, we could never have delivered this product this quickly had we had to manage the integration of servers, storage and networking required to replicate the capabilities that the ODA gave us out of the box. Future customers will get all the benefit of the lessons we learned from this first experience, and we expect future deploy and configuration costs will be negligible.”

Key DecisionHub Features

- **Management of data flows** into, out of and within the your product platform
 - ETL – bulk data movement among various databases
 - EAI – integration with message based communication among business services
- **BI - Business Intelligence**
 - Central Enterprise Data Warehouse providing unified information source
 - Specialized data marts for reporting, analytics and services like PerfectOffer
 - Sophisticated industry specific reporting
 - Support for data analysts who want to do ad-hoc reporting
- **Dashboards** – interactive and dynamic HTML dashboard server
- **Advanced Analytics** – descriptive, predictive and prescriptive analytic processing

The DW/BI *presentation area* is where data is organized, stored, and made available to users of reports, dashboards, and other analytical BI applications. Because the back room ETL system is off-limits, the presentation area is the DW/BI environment as far as the business community is concerned; it is all the business sees and touches via their access tools and BI applications.

- data be presented, stored, and accessed in dimensional schemas
- granularity down to detailed, atomic data
- structured around key business process measurement events
- all the dimensional structures built using common, conformed dimensions

DecisionHub 1.0 supports Oracle Database 12c, and is ready to run on the Oracle Database Appliance versions X4 and X5