

COMMONWEALTH OF PENNSYLVANIA

DEPARTMENT OF HUMAN SERVICES

INFORMATION TECHNOLOGY STANDARD

Name Of Standard: Operational Data Store (ODS) Development	Number: STD-EKMS010
Domain: Knowledge Management	Category: Data Warehouse
Date Issued: 2/27/2009	DSH Bureau of Information Systems
Date Revised: 2/24/2016	

Abstract:

The purpose of this Standard is to establish enterprise-wide standards and guidance for the development of subject areas within the Department of Human Services (DHS) Operational Data Store (ODS). An ODS is a “subject-oriented, integrated, current, volatile collection of data used to support the tactical decision-making process for the enterprise.”

General:

This standard applies to all DHS developers, both Commonwealth and contractor.

Standard:

Business reporting requirements that fall into the operational business intelligence report category should be fulfilled using data stored within the ODS. Report categories are summarized in the Reporting Requirements standard.

ODS Background

According to book, “Building the Operational Data Store” by Bill Inmon and Claudia Imhoff, an ODS is a “subject-oriented, integrated, current, volatile collection of data used to support the tactical decision-making process for the enterprise.”

1. Subject-oriented – The ODS is built focused around different data subjects. Examples of data subjects are customer, program, and provider data.
2. Integrated – the data from each subject area is fully integrated – much like the data in the data warehouse. It is cleansed as much as possible during the Extraction, Transformation, and Load (ETL) process before it is populated into the store.

3. Current – the data in the ODS is as current as we can technologically make it – a significant difference from a traditional data warehouse. Current as in as close to “real time” as we can get.

4. Volatile – a major differentiator from a data warehouse. The new or changed data flowing into the ODS updates the appropriate records or fields exactly like you would update a record in an OnLine Transactional Processing (OLTP) operational system. For example, when a new customer address is brought in, the old customer address is overwritten. An audit trail is created to trace the change but the old data is gone. If you want the history of the customer’s moves, you will have to either go to the audit trail or the data warehouse where history is preserved in snapshots.

It is this last aspect of the ODS that most differentiates it significantly from a data warehouse and puts it in the camp of operational systems. Because of this feature, referential integrity should be fully implemented. Cascading updates and deletes, complete edit checks, and so on are needed just as they are in operational systems.

The ODS concept was originally developed to integrate both current master data and current transactional data. The ODS’ purpose ranges from producing operational reports to propagating operational data for downstream operational systems to supporting the migration of legacy systems. And, of course, like any other operational database – it can be a source of data for the data warehouse.

ODS Standards

As operational data is changed, the ODS data get changed. No history of changes is kept. For data that is updated in the transactional system, a snapshot of only active records and records updated within the period of the EDW update frequency should be taken. This data will overlay the data currently in the ODS to maintain a single snapshot within the ODS. An ODS should not be used to delay or prevent a more frequent update of the data warehouse when requirements dictate a more frequent EDW update.

As a rule, the ODS should be centralized and integrated across the enterprise. Current master data should utilize the same data structure to ensure integration points.

A goal of the ODS is to source EDW loads from ODS data to prevent duplicate processing on the transactional system and the transformation servers and to preserve a single source of data. In order to facilitate the use of the operational data store (ODS) as a source for the enterprise data warehouse (EDW), all columns from detail transactional system tables should be brought into the ODS structure.

There may be instances where data for the EDW must be directly sourced from the transactional system. If possible, the EDW should be sourced from the ODS to minimize load and activity on the operational (OLTP) system.

Exemptions from this Standard:

There will be no exemptions to this standard.

Refresh Schedule:

All standards and referenced documentation identified in this standard will be subject to review and possible revision annually or upon request by the DHS Information Technology Standards Team.

Standard Revision Log:

Change Date	Version	Change Description	Author and Organization
11/17/2007	1.0	Initial Creation	EKMS
2/24/2016	1.1	Reviewed for content changes; changed DPW to DHS.	Don Pidich - EKMS