

BUSINESS INTELLIGENCE

The goals of a Data Warehouse

Business intelligence is a set of concepts, methods and processes to make better business decisions, using information from various sources and applying expertise and skills to develop an accurate understanding of business dynamics.¹ From the technological point of view, it is a wide category of applications and technologies for gathering, storing, accessing and analyzing data in order to help the Company's information users making better business decisions based on high-quality information.²

One of the most important assets of any organization is its information. This asset is almost always kept by an organization in two forms: the operational systems of record and the data warehouse. Crudely speaking, the operational systems are where the data is put in, and the data warehouse is where we get the data out. The data warehouse has profoundly different needs, clients, structures, and rhythms than the operational systems of record. Unfortunately, we continue to encounter supposed data warehouses that are mere copies of the operational system of record stored on a separate hardware platform.

Goals of a Data Warehouse

The data warehouse must make an organization's information easily accessible. The contents of the data warehouse must be understandable. The data must be intuitive and obvious to the business user, not merely the developer. Understandability implies legibility; the contents of the data warehouse need to be labeled meaningfully. Business users want to separate and combine the data in the warehouse in endless combinations, a process commonly referred to as slicing and dicing. The tools that access the data warehouse must be simple and easy to use. They also must return query results to the user with minimal wait times.

¹ <http://www.managementlogs.com/2004/09/what-is-bi-really-multiple-versions-of.html>

² Taken from a slide of Business Intelligence course. Profesor: Mauricio Ruiz Valdivieso. Universidad de los Andes

The data warehouse must present the organization's information consistently.

The data in the warehouse must be credible. Data must be carefully assembled from a variety of sources around the organization, cleansed, quality assured, and released only when it is fit for user consumption. Information from one business process should match with information from another. If two performance measures have the same name, then they must mean the same thing. Conversely, if two measures don't mean the same thing, then they should be labeled differently. Consistent information means high-quality information. It means that all the data is accounted for and complete. Consistency also implies that common definitions for the contents of the data warehouse are available for users.

The data warehouse must be adaptive and resilient to change.

We simply can't avoid change. User needs, business conditions, data, and technology are all subject to the shifting sands of time. The data warehouse must be designed to handle this inevitable change. Changes to the data warehouse should be graceful, meaning that they don't invalidate existing data or applications. The existing data and applications should not be changed or disrupted when the business community asks new questions or new data is added to the warehouse. If descriptive data in the warehouse is modified, we must account for the changes appropriately.

The data warehouse must be a secure bastion that protects our information assets.

An organization's informational crown jewels are stored in the data warehouse. At a minimum, the warehouse likely contains information about what we're selling to whom at what price—potentially harmful details in the hands of the wrong people. The data warehouse must effectively control access to the organization's confidential information.

The data warehouse must serve as the foundation for improved decision making.

The data warehouse must have the right data in it to support decision making. There

The goals can be developed by walking through the halls of any organization and listening to business management. Inevitably, these recurring themes emerge:

"We have mountains of data in this company, but we can't access it."

"We need to slice and dice the data every which way."

"You've got to make it easy for business people to get at the data directly."

"Just show me what is important."

"It drives me crazy to have two people present the same business metrics at a meeting, but with different numbers."

"We want people to use information to support more fact-based decision making."

These concerns are so universal that they drive the bedrock requirements for the data warehouse. Let's turn these business management quotations into data warehouse requirements.

is only one true output from a data warehouse: the decisions that are made after the data warehouse has presented its evidence. These decisions deliver the business impact and value attributable to the warehouse. The original label that predates the data warehouse is still the best description of what we are designing: a decision support system.

The business community must accept the data warehouse if it is to be deemed successful. It doesn't matter that we've built an elegant solution using best-of-breed products and platforms. If the business community has not embraced the data warehouse and continued to use it actively six months after training, then we have failed the acceptance test. Unlike an operational system rewrite, where business users have no choice but to use the new system, data warehouse usage is sometimes optional. Business user acceptance has more to do with simplicity than anything else.

As this list illustrates, successful data warehousing demands much more than being a stellar DBA or technician. With a data warehousing initiative, we have one foot in our information technology (IT) comfort zone, while our other foot is on the unfamiliar turf of business users. We must straddle the two, modifying some of our tried-and-true skills to adapt to the unique demands of data warehousing. Clearly, we need to bring a bevy of skills to the party to behave like we're a hybrid DBA/MBA.

CONFIDENTIAL

