



"Linked Open Apps Ecosystem to open up innovation in smart cities"

Project Number: 297363

Deliverable:	D4.3 Data warehouse & BI
Version:	1.6
Delivery date:	06/02/2013
Dissemination level:	PU
Author:	RETEVISION

Statement of originality:

This deliverable contains original unpublished work except where clearly indicated otherwise.

Acknowledgement of previously published material and of the work of others has been made through appropriate citation, quotation or both.

Summary:

This deliverable presents the first version of DWH & BI prototype of iCity platform. The document also analyses different types of Business intelligence products and describes differences between relational and non-relational databases.

DOCUMENT HISTORY

Version	Date of issue	Status	Content and changes	Modified by
0.1	05/10/2012	Draft	Creation	RETE
0.2	17/01/2013	Draft	Analysis of DWH & BI: <ul style="list-style-type: none"> • Existing solutions • BI software benchmarking • Open source BI software vs Commercial BI software • Barcelona BI Interface description 	BCN
0.3	28/01/2013	Draft	Analysis of DWH & BI <ul style="list-style-type: none"> • DWH key points 	CISCO
0.4	30/01/2013	Draft	iCity DWH & BI prototype description	RETE
0.5	05/02/2013	Draft	Genoa BI interface description	CDG
0.6	06/02/2013	Draft peer review	Minor changes. Final draft version peer review.	RETE

Document contributors

Partner	Contributor
Retevisión	Carmen Vicente, Javier Marcos, Luis Moreno y Elena Villa
BCN	Marc Garriga and Miriam Alvarado
CISCO	Frank Van Steenwinkel
CDG	Paolo Castiglieri

TABLE OF CONTENTS

- 1. Introduction 6**
- 1.1 Purpose 6
- 1.2 Scope..... 6
- 1.3 Overview..... 6
- 2. Analysis of DWH & BI..... 7**
- 2.1 Existing Solutions..... 7
 - 2.1.1 ETL: Extract, Transform and Load 7
 - 2.1.2 Data Warehouse..... 8
 - 2.1.3 Business Intelligence 10
- 2.2 BI Software Benchmarking 11
 - 2.2.1 Software platform capabilities 11
 - 2.2.2 Evaluation Criteria 14
 - 2.2.3 Quadrant Descriptions 16
 - 2.2.4 Major Vendor Strengths and Cautions 17
- 2.3 Open Source BI Software vs. Commercial BI Software 26
 - 2.3.1 What is open source business intelligence? 26
 - 2.3.2 Advantages of open source BI tools 26
 - 2.3.3 Disadvantages of open source BI tools..... 27
- 2.4 Description of Barcelona and Genoa BI interfaces 27
 - 2.4.1 Barcelona BI Interface 27
 - 2.4.2 Genoa BI Interface 28
- 2.5 DWH key points 28
 - 2.5.1 Relational or non-relational Data Base 28
 - 2.5.1 Video management 29
- 3. iCity DWH & BI Prototype 32**
- 3.1 DWH adaptation..... 32
- 4. Conclusions..... 35**

TABLE OF FIGURES

Figure 1: ETL	7
Figure 2: Magic Quadrant for Business Intelligence Platforms. Source: Gartner (February 2012).....	14
Figure 3: Ability to Execute Evaluation Criteria. Source: Gartner (February 2012).....	15
Figure 4: Completeness of Vision Evaluation Criteria. Source: Gartner (February 2012).....	16
Figure 5 Dedicated memcached machines.....	29
Figure 6 Index Store	29
Figure 7 Show and Share components.....	30
Figure 8 Show and Share architecture from client perspective	30
Figure 9: Data Warehouse prototype	33
Figure 10: Example of data base model	34

Abbreviations and Acronyms

Acronym	Description
DWH, DW	Data Warehouse
BI	Business Intelligence
WP	Work Package
SME	Small Medium Enterprises
ETL	Extract, Transform, Load
EDW	Enterprise Data Warehouse
ODS	Operational Data Store
DSS	Decision Support System
RBDMS	Relational Database Management System
API	Application Programming Interface
OGC	Open Geospatial Consortium
D4.x	Deliverable 4.x

1. Introduction

This document reports the work achieved in WP4 regarding Data Warehouse adaptation and Business Intelligence.

1.1 Purpose

The aim of this deliverable is to set up the first version of Data Warehouse prototype, which has the main goal of providing a homogeneous storage data, although data is coming from heterogeneous open infrastructures. Besides, Data Warehouse will allow future deployments of Business Intelligence, based on WP3 design requirements, WP5 pilots' needs and WP7 exploitation models definition.

This activity is alive until the end of the project and the DWH & BI prototype will be enriched following the inputs coming from WP3, WP5 and WP7. Thus, future versions of the deliverable will be provided: D4.8 (M24) and D4.13 (M35).

1.2 Scope

The deliverable is mainly focused on the state of the art related Data Warehouse and Business Intelligence solutions, in order to establish the implementation guidelines for both topics, covering the iCity requirements.

As a result, the iCity platform first version prototype adapts a simple Data Warehouse system which will be improved with more functionalities and features in future versions of the iCity platform prototype.

1.3 Overview

First part of the deliverable is an analysis of data warehouse and business intelligence solutions based on the state of the art. The analysis presents existing solutions and also elaborates comparisons, in order to establish the best guidelines for deploying data warehouse and BI to iCity platform.

Regarding BI, not only commercial solutions have been analyzed. Thus also open source solutions have been taken into account. Regarding DWH, as a large amount of data is expected to be managed by iCity platform, the analysis presents a discussion about the use of relational or non-relational data base and the management of video content.

Second part of the deliverable is focused on the description of the data warehouse adaptation to iCity prototype.

2. Analysis of DWH & BI

2.1 Existing Solutions

In this chapter, we should explain the following items:

- BI Analysis Software for Enterprises.
- BI Analysis Software for SME.
- Analysis of advantages and disadvantage of implantation an Open Source vs. Commercial Product.

The main concepts in order to explain BI existing solutions are listed below¹.

2.2 ETL: Extract, Transform and Load

Extract, Transform and Load (ETL) refers to a process in database usage and especially in data warehousing that involves:

- 1) Extracting data from outside sources
- 2) Transforming it to fit operational needs (which can include quality levels)
- 3) Loading it into the end target (database, more specifically).

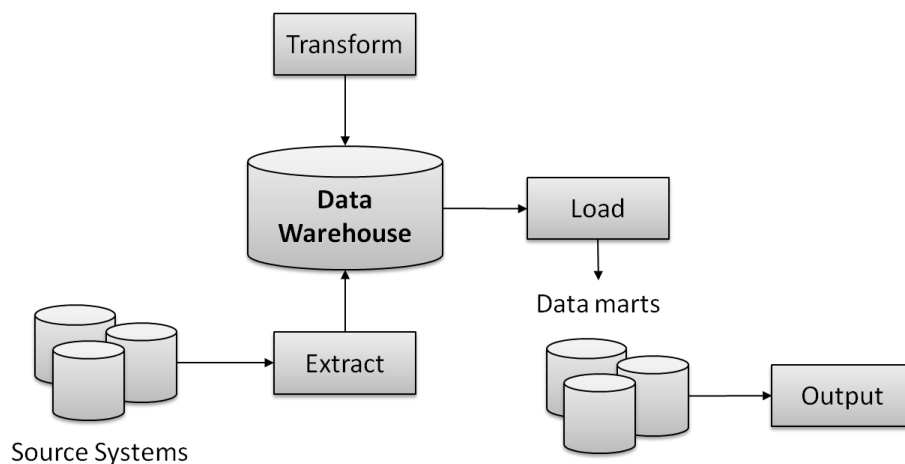


Figure 1: ETL

Extract

The first part of an ETL process involves extracting the data from the source systems. In many cases this is the most challenging aspect of ETL because data is from different source systems, different formats, different structures, different storages, etc., in fact extracting data correctly will set the stage for how subsequent processes will go.

¹ Most of the information in this section is taken from Wikipedia.

Transform

The transform stage applies a series of rules or functions to the extracted data from the source to derive the data for loading into the end target.

Load

The load phase loads the data into the end target, usually the data warehouse (DW). Depending on the requirements of the organization, this process varies widely. Some data warehouses may overwrite existing information with cumulative information, frequently updating extract data is done on daily, weekly or monthly basis. Other DW (or even other parts of the same DW) may add new data in a historical form, for example, hourly. To understand this, consider a DW that is required to maintain sales records of the last year. Then, the DW will overwrite any data that is older than a year with newer data. However, the entry of data for any one year window will be made in a historical manner. The timing and scope to replace or append are strategic design choices dependent on the time available and the business needs. More complex systems can maintain a history and audit trail of all changes to the data loaded in the DW.

As the load phase interacts with a database, the constraints defined in the database schema — as well as in triggers activated upon data load — apply (for example, uniqueness, referential integrity, mandatory fields), which also contribute to the overall data quality performance of the ETL process.

For example, a financial institution might have information on a customer in several departments and each department might have that customer's information listed in a different way. The membership department might list the customer by name, whereas the accounting department might list the customer by number. ETL can bundle all this data and consolidate it into a uniform presentation, such as for storing in a database or data warehouse.

Another way that companies use ETL is to move information to another application permanently. For instance, the new application might use another database vendor and most likely a very different database schema. ETL can be used to transform the data into a format suitable for the new application to use.

2.3 Data Warehouse

In computing, a data warehouse or enterprise data warehouse (DW, DWH, or EDW) is a database used for reporting and data analysis. It is a central repository of data which is created by integrating data from multiple disparate sources. Data warehouses store current as well as historical data and are used for creating trending reports for senior management reporting such as annual and quarterly comparisons.

The data stored in the warehouse are uploaded from the operational systems (such as marketing, sales etc., shown in the figure to the right). The data may pass through an operational data store for additional operations before they are used in the DW for reporting.

The typical ETL-based data warehouse uses staging, data integration, and access layers to house its key functions. The staging layer or staging database stores raw data extracted from each of the disparate source data systems. The integration layer integrates the disparate data sets by transforming the data from the staging layer often storing this transformed data in an operational data store (ODS) database. The integrated data are then moved to yet another database, often called the data warehouse database, where the data is arranged into hierarchical groups often called dimensions and into facts and aggregate facts. The combination of facts and dimensions is sometimes called a star schema. The access layer helps users retrieve data.

A data warehouse constructed from an integrated data source system does not require ETL, staging databases, or operational data store databases. The integrated data source systems may be considered to be a part of a distributed operational data store layer. Data federation methods or data virtualization methods may be used to access the distributed integrated source data systems to consolidate and aggregate data directly into the data warehouse database tables. Unlike the ETL-based data warehouse, the integrated source data systems and the data warehouse are all integrated since there is no transformation of dimensional or reference data. This integrated data warehouse architecture supports the drill down from the aggregate data of the data warehouse to the transactional data of the integrated source data systems.

Data warehouses can be subdivided into data marts. Data marts store subsets of data from a warehouse.

This definition of the data warehouse focuses on data storage. The main source of the data is cleaned, transformed, catalogued and made available for use by managers and other business professionals for data mining, online analytical processing, market research and decision support (Marakas & O'Brien 2009). However, the means to retrieve and analyze data, to extract, transform and load data, and to manage the data dictionary are also considered essential components of a data warehousing system. Many references to data warehousing use this broader context. Thus, an expanded definition for data warehousing includes business intelligence tools, tools to extract, transform and load data into the repository, and tools to manage and retrieve metadata.

Benefits of a data warehouse

A data warehouse maintains a copy of information from the source transaction systems. This architectural complexity provides the opportunity to:

- Maintain data history, even if the source transaction systems do not.
- Integrate data from multiple source systems, enabling a central view across the enterprise. This benefit is always valuable, but particularly so when the organization has grown by merger.
- Improve data quality, by providing consistent codes and descriptions, flagging or even fixing bad data.
- Present the organization's information consistently.
- Provide a single common data model for all data of interest regardless of the data's source.
- Restructure the data so that it makes sense to the business users.
- Restructure the data so that it delivers excellent query performance, even for complex

analytic queries, without impacting the operational systems.

- Add value to operational business applications, notably customer relationship management (CRM) systems.

2.4 Business Intelligence

Business intelligence (BI) is the ability of an organization to collect, maintain, and organize data. This produces large amounts of information that can help develop new opportunities. Identifying these opportunities, and implementing an effective strategy, can provide a competitive market advantage and long-term stability.

BI technologies provide historical, current and predictive views of business operations. Common functions of business intelligence technologies are reporting, online analytical processing, analytics, data mining, process mining, complex event processing, business performance management, benchmarking, text mining, predictive analytics and prescriptive analytics.

The goal of modern business intelligence deployments is to support better business decision-making. Thus a BI system can be called a decision support system (DSS). Though the term business intelligence is sometimes a synonym for competitive intelligence (because they both support decision making), BI uses technologies, processes, and applications to analyse mostly internal, structured data and business processes while competitive intelligence gathers, analyses and disseminates information with a topical focus on company competitors. If understood broadly, business intelligence can include the subset of competitive intelligence.

Business intelligence and data warehousing

Often BI applications use data gathered from a data warehouse or a data mart. However, not all data warehouses are used for business intelligence, nor do all business intelligence applications require a data warehouse.

To distinguish between the concepts of business intelligence and data warehouses, Forrester Research often defines business intelligence in one of two ways:

Using a broad definition: "Business Intelligence is a set of methodologies, processes, architectures, and technologies that transform raw data into meaningful and useful information used to enable more effective strategic, tactical, and operational insights and decision-making". When using this definition, business intelligence also includes technologies such as data integration, data quality, data warehousing, master data management, text and content analytics, and many others that the market sometimes lumps into the Information Management segment. Therefore, Forrester refers to data preparation and data usage as two separate, but closely linked segments of the business intelligence architectural stack.

Forrester defines the latter, narrower business intelligence market as, "...referring to just the top layers of the BI architectural stack such as reporting, analytics and dashboards."

2.5 BI Software Benchmarking

An extended analysis on Business Intelligence is The Magic Quadrant for Business Intelligence Platforms² published by Gartner.

This study presents a global view of Gartner's opinion of the main software vendors that should be considered by organizations seeking to develop business intelligence (BI) applications. To be included in the Magic Quadrant, vendors must generate at least \$15 million in BI-related software license revenue annually. The Magic Quadrant is based on a customer survey including vendor-provided references, as well as survey responses from BI users from Gartner's BI Summit. On 2011 there were 1,364 survey responses.

2.6 Software platform capabilities

As defined by Gartner, business intelligence (BI) platforms enable all types of users — from IT staff to consultants to business users — to build applications that help organizations learn about and understand their business. Gartner defines a BI platform as a software platform that delivers the 14 capabilities listed below. These capabilities are organized into three categories of functionality: integration, information delivery and analysis. Information delivery is the core focus of most BI projects today, but there is an increased interest in deployments of analysis to discover new insights, and in integration to implement those insights.

2.6.1.1 Integration

BI infrastructure — All tools in the platform use the same security, metadata, administration, portal integration, object model and query engine, and should share the same look and feel.

Metadata management — Not only should all tools leverage the same metadata, but the offering should provide a robust way to search, capture, store, reuse and publish metadata objects such as dimensions, hierarchies, measures, performance metrics and report layout objects.

Development tools — The BI platform should provide a set of programmatic development tools and a visual development environment, coupled with a software developer's kit for creating BI applications, integrating them into a business process, and/or embedding them in another application. The BI platform should also enable developers to build BI applications without coding by using wizard-like components for a graphical assembly process. The development environment should also support Web services in performing common tasks such as scheduling, delivering, administering and managing. In addition, the BI application can assign and track events or tasks allotted to specific users, based on predefined business rules. Often, this capability can be delivered by integrating with a separate portal or workflow tool.

Collaboration — This capability enables BI users to share and discuss information, BI content and results, and/or manage hierarchies and metrics via discussion threads, chat and annotations, either embedded in the BI platform or through integration with collaboration, social software and analytical master data management (MDM).

² Published: 6 February 2012 ID:G00225500

Analyst(s): John Hagerty, Rita L. Sallam, James Richardson

<http://businessintelligence.info/docs/estudios/Magic-Quadrant-for-Business-Intelligence-Platforms-2012.pdf>

2.6.1.2 Information Delivery

Reporting — Reporting provides the ability to create formatted and interactive reports, with or without parameters, with highly scalable distribution and scheduling capabilities. In addition, BI platform vendors should handle a wide array of reporting styles (for example, financial, operational and performance dashboards), and should enable users to access and fully interact with BI content delivered consistently across delivery platforms including the Web, mobile devices and common portal environments.

Dashboards — This subset of reporting includes the ability to publish formal, Web-based or mobile reports with intuitive interactive displays of information, including dials, gauges, sliders, check boxes and traffic lights. These displays indicate the state of the performance metric compared with a goal or target value. Increasingly, dashboards are used to disseminate real-time data from operational applications or in conjunction with a complex event processing engine.

Ad hoc query — This capability enables users to ask their own questions of the data, without relying on IT to create a report. In particular, the tools must have a robust semantic layer to allow users to navigate available data sources. These tools should include a disconnected analysis capability that enables users to access BI content and analyse data remotely without being connected to a server-based BI application. In addition, these tools should offer query governance and auditing capabilities to ensure that queries perform well.

Microsoft Office integration — In some use cases, BI platforms are used as a middle tier to manage, secure and execute BI tasks, but Microsoft Office (particularly Excel) acts as the BI client. In these cases, it is vital that the BI vendor provides integration with Microsoft Office applications, including support for document and presentation formats, formulas, data "refreshes" and pivot tables. Advanced integration includes cell locking and write-back.

Search-based BI — This applies a search index to both structured and unstructured data sources and maps them into a classification structure of dimensions and measures (often, but not necessarily leveraging the BI semantic layer) that users can easily navigate and explore using a search (Google-like) interface. This capability extends beyond keyword searching of BI platform content and metadata.

Mobile BI — This capability enables organizations to deliver report and dashboard content to mobile devices (such as smartphones and tablets) in a publishing and/or interactive (bidirectional) mode, and takes advantage of the interaction mode of the device (tapping, swiping and so on) and other capabilities not commonly available on desktops and laptops, such as location awareness.

2.6.1.3 Analysis

Online analytical processing (OLAP) — This enables end users to analyze data with extremely fast query and calculation performance, enabling a style of analysis known as "slicing and dicing." Users are (often) able to easily navigate multidimensional drill paths. And they (sometimes) have the ability to write-back values to a proprietary database for planning and "what if" modeling purposes. This capability could span a variety of data architectures (such as relational or multidimensional) and storage architectures (such as disk-based or in-memory).

Interactive visualization — This gives users the ability to display numerous aspects of the data more efficiently by using interactive pictures and charts, instead of rows and columns. Over time, advanced visualization will go beyond just slicing and dicing data to include more process-driven BI projects, allowing all stakeholders to better understand the workflow through a visual representation.

Predictive modeling and data mining — This capability enables organizations to classify categorical variables and to estimate continuous variables using advanced mathematical techniques. BI developers are able to integrate models easily into BI reports, dashboards and analysis, and business processes.

Scorecards — These take the metrics displayed in a dashboard a step further by applying them to a strategy map that aligns key performance indicators (KPIs) with a strategic objective. Scorecard metrics should be linked to related reports and information in order to do further analysis. A scorecard implies the use of a performance management methodology such as Six Sigma or a balanced scorecard framework.

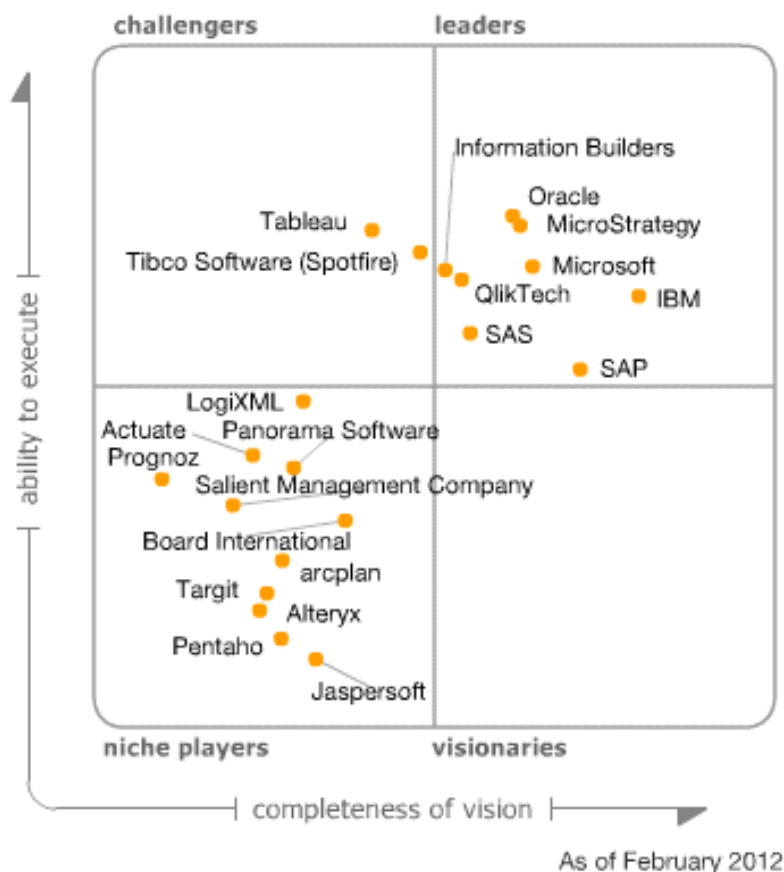


Figure 2: Magic Quadrant for Business Intelligence Platforms. Source: Gartner (February 2012)

2.7 Evaluation Criteria

2.7.1.1 Ability to Execute

Vendors are judged on their ability and success in making their vision a market reality. In addition to the opinions of Gartner's analysts, the scores are based on three sources: customer perceptions of each vendor's strengths and challenges derived from BI-related inquiries with Gartner; an online survey of vendor customers conducted in late 2011, yielding 1,364 responses; and a vendor-completed questionnaire about the vendor's BI strategy and operations.

- **Product/Service:** How competitive and successful are the goods and services offered by the vendor in this market? This includes current product/service capabilities, quality, feature sets and skills, whether offered natively or through OEM agreements/partnerships.
- **Overall Viability:** What is the likelihood of the vendor continuing to invest in products and services for its customers? Viability includes an assessment of the overall organization's financial health, the financial and practical success of the business unit, and the likelihood of the individual business unit to continue to invest in the product, continue to offer the product and advance the state of the art within the organization's portfolio of products.

- **Sales Execution/Pricing:** Does the vendor provide cost-effective licensing and maintenance options? This covers the technology provider's capabilities in all presales activities and the structure that supports them. This includes deal management, pricing and negotiation, presales support and the overall effectiveness of the sales channel.
- **Market Responsiveness and Track Record:** Can the vendor respond to changes in market direction as customer requirements evolve? This covers the ability to respond, change direction, be flexible and achieve competitive success as opportunities develop, competitors act, customer needs evolve and market dynamics change. This criterion also considers the provider's history of responsiveness.
- **Marketing Execution:** Are customers aware of the vendor's offerings in the market? This assesses the clarity, quality, creativity and efficacy of programs designed to deliver the organization's message in order to influence the market, promote the brand and business, increase awareness of the products and establish a positive identification with the product/brand and organization in the minds of buyers. This mind share can be driven by a combination of publicity, promotional, thought leadership, word-of-mouth and sales activities. This criterion was not rated separately this year and therefore was given a "no rating" in the Magic Quadrant model. Instead, our assessment of Market Execution was combined with Market Responsiveness and Track Record into one criterion on this year's Magic Quadrant.
- **Customer Experience:** How well does the vendor support its customers? How trouble-free is the software?
- **Operations:** What is the ability of the organization to meet its goals and commitments? This criterion was given a "no rating." Assessment of a vendor's ability to meet its goals and commitments is incorporated into the Market Responsiveness and Track Record criterion.

<u>Evaluation Criteria</u>	<u>Weighting</u>
Product/Service	high
Overall Viability (Business Unit, Financial, Strategy, Organization)	high
Sales Execution/Pricing	high
Market Responsiveness and Track Record	standard
Marketing Execution	no rating
Customer Experience	high
Operations	no rating

Figure 3: Ability to Execute Evaluation Criteria. Source: Gartner (February 2012)

2.7.1.2 Completeness of Vision

Vendors are rated on their understanding of how market forces can be exploited to create value for customers and opportunity for themselves. The scores are based on the same sources like the *Ability to Execute* criteria.

- **Market Understanding:** Does the vendor have the ability to understand buyers' needs, and to translate those needs into products and services?
- **Marketing Strategy:** Does the vendor have a clear set of messages that communicate its value and differentiation in the market?

- **Sales Strategy:** Does the vendor have the right combination of direct and indirect resources to extend its market reach?
- **Offering (Product) Strategy:** Does the vendor's approach to product development and delivery emphasize differentiation and functionality that maps to current and future requirements? The major business analytics market growth drivers described in the Market Overview section of this report were used as a rubric to assess both the Offering (Product) Strategy and Innovation criteria, which are combined into one score this year.
- **Business Model:** How sound and logical is the vendor's underlying business proposition? Note that this criterion has been given a "no rating" because all vendors in the market have a viable business model.
- **Vertical/Industry Strategy:** How well can the vendor meet the needs of various industries, such as financial services or the retail industry?
- **Innovation:** How well does the vendor direct related, complementary and synergistic layouts of resources, expertise or capital for investment, consolidation, defensive or pre-emptive purposes? How well does the vendor exploit current or new technologies and combine them in a novel way to address a market need? Innovation and Offering (Product) Strategy are combined into one score for the purpose of this year's Magic Quadrant.
- **Geographic Strategy:** How well can the vendor meet the needs of locations outside its native country, either directly or through partners?

Evaluation Criteria	Weighting
Market Understanding	high
Marketing Strategy	high
Sales Strategy	high
Offering (Product) Strategy	high
Business Model	no rating
Vertical/Industry Strategy	standard
Innovation	no rating
Geographic Strategy	standard

Figure 4: Completeness of Vision Evaluation Criteria. Source: Gartner (February 2012)

2.7.2 Quadrant Descriptions

2.7.2.1 Leaders

Leaders are vendors that are reasonably strong in the breadth and depth of their BI platform capabilities and can deliver on enterprise wide implementations that support a broad BI strategy. Leaders articulate a business proposition that resonates with buyers, supported by the viability and operational capability to deliver on a global basis.

2.7.2.1 Challengers

Challengers offer a good breadth of BI platform functionality and are well positioned to succeed in the market. However, they may be limited to specific use cases, technical environments or application domains. Their vision may be hampered by a lack of coordinated strategy across the various products in their BI platform portfolio, or they may lack the

marketing effort, sales channel, geographic presence, industry-specific content, and awareness offered by the vendors in the Leaders quadrant.

2.7.2.2 Visionaries

Visionaries are vendors that have a strong vision for delivering a BI platform. They are distinguished by the openness and flexibility of their application architectures, and they offer depth of functionality in the areas they address, but they may have gaps relating to broader functionality requirements. A Visionary is a market thought-leader and innovator. However, it may have yet to achieve sufficient scale — or there may be concerns about its ability to grow and provide consistent execution.

2.7.2.3 Niche Players

Niche Players are those that do well in a specific segment of the BI platform market — such as reporting or dashboarding — or that have limited capability to innovate or outperform other vendors in the market. They may focus on a specific domain or aspect of BI, but are likely to lack depth of functionality elsewhere. Or they may have gaps relating to broader BI platform functionality. Alternatively, Niche Players may have a reasonably broad BI platform, but have limited implementation and support capabilities or relatively limited customer bases, such as in a specific geography or industry. Or they may not yet have achieved the necessary scale to solidify their market positions.

2.8 Major Vendor Strengths and Cautions

We have considered the principal vendors to analyse as:

- IBM
- Oracle
- SAP
- Microsoft
- Pentaho
- SAS
- MicroStrategy
- QlikTech
- Tableau

The conclusions of the Gartner analysis about these vendors are:

IBM

Strengths

- ✓ IBM maintains its leading position on the Completeness of Vision axis for 2012 Magic Quadrant. The company takes a holistic approach to what it calls Business Analytics and Optimization (BAO), combining comprehensive software, hardware and services in a coordinated market offering. IBM's business analytics software portfolio includes a unified BI, analytics and performance management platform, and is complemented by IBM information management software and appliances.
- ✓ In 4Q10, IBM introduced its latest business analytics platform, IBM Cognos 10. Throughout 2011, additional capabilities have been released and customer adoption has begun in earnest.

- ✓ Advanced analytics is a particular IBM strength. The company's SPSS software continues to advance nicely, readily allowing IBM to bid for predictive analytics and statistical use cases.
- ✓ The top reasons why customers select IBM are functionality, ease of use for end users, and data access and integration. IBM's road map and future vision weighed heavily in reference decisions. In 2011, IBM delivered a new Cognos 10 mobile application for the iPad that is included free in existing user roles. In early 2012 the company will introduce Cognos Insight, a personal, desktop BI product that enables independent discovery and "what if" modeling, while also providing full interoperability with the larger workgroup and enterprise solutions.

Cautions

- Again this year, references consider the Cognos products more difficult to implement and use than those of competitors. References indicate that Cognos software is used largely by a consumer/casual user population. Reporting is the most extensively deployed component, followed by ad hoc query and OLAP analysis.
- IBM's customers also continue to have less than optimal customer experiences, with support and sales interactions, along with product quality..
- License cost continues to be another source of customer concern across all products in the IBM business analytics portfolio. Higher than expected costs to upgrade from Cognos 8 to Cognos 10 have stalled some projects, but changes in configuration, user roles, and/or support costs appear to drive the increase. As a counterpoint, existing Cognos 10 users did not identify license cost as a concern.

Microsoft

Strengths

- ✓ Microsoft offers a competitive set of BI capabilities, packaging and pricing that appeal to Microsoft developers and its independent distributor channel. The company has consistently invested in building and enhancing BI capabilities into three of its core offerings — Microsoft Office (specifically Excel), Microsoft SQL Server and Microsoft SharePoint — in order to increase their value and drive upgrades. By incorporating BI capabilities into its most ubiquitous products, Microsoft virtually guarantees its BI offering's continued adoption, particularly in organizations with a Microsoft-centric information infrastructure. As a result of this strategy, since the company's serious entry into the market in 2000, Microsoft's BI market share has grown steadily to take the No. 3 spot in 2010.
- ✓ Microsoft's low-license-cost bundling strategy for BI platforms makes it a compelling license-cost value proposition for organizations that want to deploy BI to a wider range of users, or that want to lower overall BI portfolio license costs by using lower-cost BI tools for basic BI functions. Its license cost profile is comparable to open-source BI vendors, and is considerably less than its commercial competitors.
- ✓ Microsoft's market success is also driven in part by its IT-oriented, BI authoring tools within SQL Server, which are based on Visual Studio, the broadly adopted development environment. This approach, along with targeted marketing efforts and

programs for building strong developer communities and support, has helped Microsoft lower the cost and expand the availability of its BI skills.

- ✓ While Microsoft has traditionally focused on the developer, it continues to enhance reporting, dashboarding and data discovery capabilities in Excel with the intention of making Excel not only the most widely deployed BI tool, but also the most functional for business users.
- ✓ Use of OLAP functionality by Microsoft customers is among the highest when compared to other vendors. This can be attributed to the success and adoption of Microsoft SQL Server Analysis Services functionality bundled with Microsoft SQL Server and its optimizations with Microsoft front-end tools.
- ✓ Microsoft's cloud-based DataMarket offering, which makes external data easier to consume, analyse and integrate with internal data, is a unique enhancement to Microsoft's portfolio of BI capabilities. DataMarket is an online data market that enables ISVs and business users to access, purchase and analyse trusted, public-domain and commercial premium data. ISVs can use this data to build new analytic applications. Business users can incorporate and analyse this external data with internal data sources using Microsoft Excel and PowerPivot, or with partner tools, such as those from Tableau Software.

Cautions

- Since Gartner began surveying BI platform customers for this Magic Quadrant research five years ago, this is the first year that Microsoft has scored below the survey average on key Ability to Execute measures, including overall product functionality, support and customer experience.
- Multiproduct complexity is a challenge. Because Microsoft's BI platform capabilities exist across three different tools (Office, SQL Server and SharePoint) that also perform non-BI functions, integrating the necessary components and building the applications is left to the organization. Microsoft's do-it-yourself approach puts more of the BI solutions development and integration onus for the platform components on customers, compared with the all-in-one purpose-built BI platforms offered by most other vendors in the BI market. Microsoft's road map for Office, which features the consolidation of more and more front-end reporting, dashboard and analysis capabilities in Excel, should begin to address some of this complexity over time.
- Microsoft lags behind most other BI vendors in delivering mobile BI capabilities. It has, instead, relied on partners to build mobile solutions for Apple iOS that integrate with Microsoft BI components. Microsoft BI assets can run in a browser today, but they are not optimized for iOS, Android or Windows devices.
- There is currently no single business metadata layer or capability that spans Microsoft's BI platform components, and there are limited capabilities for sophisticated metadata modeling, impact analysis, data lineage and change management.

MicroStrategy

Strengths

- ✓ MicroStrategy specializes in enterprise BI deployments running on top of large enterprise data warehouses. Its customers cite functionality, performance and support for large data volumes as top reasons for selecting it as a vendor. Its deployments are among the most complex in terms of large numbers of users, the highest data volume, broad product functionality use, wide deployment across an enterprise, and complexity of analytic workload, and its customers have a high level of satisfaction with product functionality.
- ✓ MicroStrategy has a focused vision that maps to key high-value market requirements, particularly for mobility, and large and diverse data, including social media data sources. The company was one of the first vendors to invest heavily in deploying BI applications on mobile devices, with earlier successes than its competitors in accumulating a respectable number of large production mobile deployments. Free trials and online training make it easy for developers to try and succeed with mobile development. Beyond mobility, MicroStrategy continues to reinforce its enterprise-scale pedigree through initiatives for high performance across all layers of its platform and against extremely large and diverse datasets. MicroStrategy has invested heavily in creating a cloud offering that includes its platform and complementary technologies, including ETL and data warehousing. Social data is another forward-looking area of focus for MicroStrategy. This past year, the company delivered a Facebook connector to enable organizations to integrate Facebook profile data, with user permissions, into a MicroStrategy analytic application.
- ✓ Developer productivity for building complex analytic applications is another of MicroStrategy's strengths. Its efficient, parameterized report development paradigm and object-oriented report development environment support centralized management, in which a small number of administrators can support big BI projects with many users, complex reporting and analysis requirements, and a large amount of data. With an extensive library of prebuilt objects, including metrics, prompts, filters and statistical functions, developers can create reports and other analytic content with high degrees of formatting and analytical sophistication, but with less effort and cost than many other platforms.
- ✓ In March 2011, MicroStrategy introduced a data discovery capability, Visual Insight, that complements and fully integrates with its enterprise, report-centric architecture. Visual Insight is available as a feature of Report Services reducing the need for most customers to purchase stand-alone interactive visualization/data discovery products. Visual Insight is also available in a free personal cloud-based version.
- ✓ MicroStrategy has built its BI platform from the ground up through completely organic development. The high level of integration of the individual platform components and the reusability of MicroStrategy's well architected and object-oriented semantic layer are the result of this strategy.

Cautions

- While the MicroStrategy development environment is robust and flexible, there is a steep learning curve, even for seasoned report developers building any level of analytic complexity into parameterized reports that simulate ad hoc analysis and interactive dashboards for business users. The need for interactivity beyond

parameterized reports and dashboards will only increase with broader mobile BI application user adoption.

- Even though MicroStrategy has comparatively moderate administration costs per user compared to its competitors, its customers report above average license and implementation costs per user.
- While MicroStrategy Mobile, its new social data capabilities and its personal cloud offering will increase its appeal to business users and line of business owners, the company currently sells predominantly to IT, which has a stack-centric buying tendency.

Oracle

Strengths

- ✓ In 2011, Oracle Business Intelligence Foundation Suite, with its principal component Oracle Business Intelligence Enterprise Edition (OBIEE), continued to execute on its stated top-to-bottom BI vision. This year, the products have the highest aggregate *Ability to Execute* scores.
- ✓ References select Oracle primarily for functionality, enterprise application integration, and data access capabilities. Additionally, customers indicated that they valued the products' ability to support large numbers of users. Like other megavendors, the product road map plays an important role in the evaluation process. Ease of use and cost do not factor significantly into the selection process.
- ✓ Oracle Business Intelligence Applications (OBIA) are predefined analytic applications for horizontal business processes such as finance, procurement and sales analysis. Additionally, the company also delivers vertical-specific analytic data models for industries such as retail and financial services for IT buyers looking to establish a common data model standard as the foundation for analytics.

Cautions

- References rate OBIEE as difficult to implement.
- Product functionality evaluation scores remain below average again this year, a trend that appeared in last year's report. Additionally, customer support and product quality issues are rated below the average (in the fourth and third quartiles respectively) for all vendors in this report. In fact, both support and product quality were also noted as issues that blocked further deployments within customer organizations.
- Oracle customers use the product mostly for static report viewing, parameterized reporting and scorecard capabilities, leading to below average user complexity ratings.

Pentaho

Strengths

- ✓ Pentaho makes its debut on the Magic Quadrant this year. It provides a comprehensive open-source BI platform composed of ETL, OLAP, reporting, dashboards, ad hoc analysis and data mining components, all managed from a

central BI server deployed either on-premises or in the cloud, with end-user access via the Web or mobile devices such as the iPad.

- ✓ Low license cost is central to Pentaho's value proposition. The No. 1 reason that customers choose Pentaho is for its perceived low license cost and TCO.
- ✓ Pentaho's lightweight footprint, in which the complete platform can be deployed in a small environment on a laptop, or can be integrated into an existing scalable architecture such as a grid for much larger deployments, makes it very flexible in meeting a broad range of deployment requirements. Moreover, Pentaho is an embeddable platform, making it very attractive to ISVs and internal IT shops for embedded use cases to deploy both on-premises and in the cloud.
- ✓ While open-source vendors, including Pentaho, tend to invest more heavily to achieve feature parity with the core BI functionality of commercial competitors rather than in innovation, Pentaho does have focused areas of forward-looking investment, on which it has been able to deliver quickly.

Cautions

- Pentaho's below average aggregate product scores (with the exception of predictive modelling) are still an indication that functional gaps in the platform remain. Moreover, Pentaho needs to continue to improve on both its business user tools, to meet growing requirements for intuitive and interactive analysis, and the usability and efficiency of its developer-oriented tools. Moreover, ease of use goes hand in hand with the effort to develop content. Despite the perception of low TCO as a primary reason for purchasing Pentaho, users report among the longest BI content development times of all vendors in the Magic Quadrant survey.
- Given that Pentaho's subscription-based model hinges on providing superior support, Pentaho's below average product support scores (particularly related to level of expertise) are a concern, especially since the company's product is also rated below average in terms of product quality, which tends to result in more customer interaction with support.
- Although Pentaho claims a single unified platform managed from a single server, the repositories and authoring environments remain separate, with migration to a single repository in process.

QlikTech

Strengths

- ✓ QlikTech is a marketing juggernaut; it has brand recognition many times more prominent than a firm with its current market share would expect.
- ✓ QlikTech's QlikView product is a self-contained BI platform, based on a wholly in-memory data store, with a set of well integrated BI tools.
- ✓ Gartner frequently sees companies deploy QlikView for prototyping and requirements gathering, leveraging its flexibility to engage end users, usually alongside a more traditionally modeled BI platform.

- ✓ QlikTech's customers report strong delivery of business benefits, particularly in making better information available to more users and expanding the type of analysis undertaken.
- ✓ Customers' rating of QlikView's functionality is very positive in nine out of 14 functional capabilities: dashboards, interactive visualization, mobile BI, search-based BI, scorecards, ad hoc query, Microsoft Office integration, OLAP, and development tools.

Cautions

- QlikTech's growing pains are more evident. For the first time, QlikTech's customers reported having a poor overall customer experience, and below average ratings for product quality and support.
- Gartner continues to hear rumblings of discontent from QlikTech customers about the structure of its pricing model and its high license cost.
- QlikTech faces increasing competition from larger BI vendors offering in-memory offerings and interactive visualization (particularly Microsoft SQL Server PowerPivot/Power View and MicroStrategy Visual Insight), all of which are intent on narrowing QlikView's opportunities for expansion by offering cheaper alternatives.
- QlikTech offers limited metadata management. Filling this gap requires additional cost and effort in the management of metadata to lockdown common definitions and calculations, and to conform dimensions for cross-functional analysis across QlikView applications.
- Although quick to develop simple or moderately complex dashboards, when it comes to building large, complex reports from various data sources, involving detailed logic or calculations, QlikView users reported the second slowest turnaround.

SAP

Strengths

- ✓ The combination of SAP BusinessObjects and SAP NetWeaver BW revenue accounts for the largest share of the BI platform market, with both SAP platforms continuing to support large enterprise deployments (more than twice the average for both data size and number of users). Similarly, a higher percentage of SAP cite "corporate standards" and "integration with enterprise applications" as among the top reasons why they chose SAP for BI.
- ✓ SAP has one of the largest global direct sales, support, and channel and services ecosystems. Moreover, the combination of SAP and BusinessObjects constitutes the largest installed base in the BI platforms market, which represents a significant and captive cross-sell and upsell market opportunity for SAP.
- ✓ SAP has a compelling and comprehensive product vision that addresses many key future trends including mobile, collaborative analytics, and analytics on big data. SAP complements its BI platform with forward-looking capabilities in the areas of collaboration and decision support (with its StreamWork product), text analysis integrated with its enterprise information management products, and search-based data exploration with its SAP BusinessObjects Explorer product.

- ✓ SAP is investing in industry- and domain-specific packaged applications built with SAP BusinessObjects that include a data model, ETL and business content.

Cautions

- Migration, implementation and integration choices can be confusing.
- While SAP's customers tend to have very large and global deployments, poor performance is mentioned as a problem limiting broader deployment.
- At the end of August 2011, SAP implemented its third license model change (Concurrent Session-Based Licenses [CSBLs] and Named User licenses) for SAP BusinessObjects since the Business Objects acquisition in January 2008. While there are many advantages for users in using CSBL, changing license models have contributed to confusion, and concern. Also, some customers are charged for upgrades when they expected to be provided with product at no/low cost.

SAS

Strengths

- ✓ SAS gets high marks for its global footprint and broad industry initiatives. Unlike some other BI platform vendors, SAS focuses on advanced analytical techniques, such as data mining and predictive modelling, where references acknowledge it as a leader of the pack. SAS's clients also have above average complexity scores (for the depth of use of different BI use cases) on larger than average data sources. SAS customers also access and interpret unstructured internal and external data more often than any other vendor's clients surveyed for this Magic Quadrant.
- ✓ SAS's solution-oriented analytic application approach to the market is a differentiator, giving the company the advantage of having a wide variety of cross-functional and vertically specific analytic applications out of the box for a variety of industries, including financial services, life sciences and manufacturing. While others are also adopting this approach, SAS remains in the lead. Customers also report an above average sales experience.
- ✓ The primary drivers for customers choosing SAS remain functionality and data integration. In addition, references reported that they select SAS because of availability of skills.
- ✓ On the software partnership front, SAS has partnered with a number of database vendors to push the execution of its models directly into the database management system without moving the data. Not only does this reduce data duplication and movement, it also allows SAS users to leverage the power and scalability features of the database to run predictive models against very large datasets with high performance.

Cautions

- References report that SAS is very difficult to implement and companies also indicate that the product is considered difficult to use for business users.
- SAS's dominance in predictive analytics and statistics continues to be challenged on many fronts.

- Customer references report that cost is the most common factor blocking further adoption.
- Despite SAS's success and awareness as a leader in the predictive analytics space, the company is still challenged to make it onto BI platform shortlist evaluations when predictive analytics is not a primary business requirement.

Tableau

Strengths

- ✓ For the third year in a row, Tableau is the "sweetheart" of the Magic Quadrant, with customers even more enamoured with it this year than in the last two. It gained overwhelmingly positive customer survey feedback across the board in all measures in the survey, including ease of use, functionality, product quality, product performance, support, customer relationship, success, achievement of business benefits and view of the vendor's future. These stellar results in part contributed to Tableau's strong *Ability to Execute* position, despite its relatively small size.
- ✓ Tableau is one of a number of stand-alone BI vendors delivering strong interactive visualization for analysis, dashboards, information delivery and managed analytic applications. Tableau's strong performance, even with an increasingly crowded competitive landscape, is evidence of its ability to meet the increased market demand for easy-to-use and intuitive interactive BI tools that are easy to deploy without IT assistance.
- ✓ Tableau's self-contained BI platform provides purpose-built, business-oriented data mashup ETL capabilities with data connectors that leverage Tableau's own VizQL technology (drag-and-drop operations in Tableau create a query in VizQL, which interprets and packages an SQL or MDX query to the database and then expresses the response graphically). Its columnar, in-memory data engine, which can be used as an alternative to its direct query access, enables fast performance on large and multisource datasets and on complex queries, such as very large multidimensional filters or complex co-occurrence or multipass queries. Zero programming data mashup capability, combined with an in-memory database, allows users to blend and visually analyse large amounts of diverse datasets with auto-detect relationships between multiple sources (of any format). This allows users to connect to any data source and produce a series of interactive dashboards, and highlight and visually filter and pass parameters directly from a graphic; or use filters (for example, check boxes, sliders, relative date filters and drop-down menus); or build in geographic intelligence to analyse their data. Interactive analysis can be shared with a report consumer equipped with a Web browser. The combination of exceptional ease of use with the ability to conduct sophisticated analysis, is a key reason users are exuberant with the platform.

Cautions

- Tableau's product functionality is more narrowly defined around analysis and interactive visualization. It lacks broader BI platform capabilities, such as production reporting and predictive analytics. Tableau has introduced a shareable semantic layer — a key enterprise feature — in its 7.0 release.

- Although Tableau's user counts remain below the survey average (albeit growing from last year), it is still largely departmentally deployed with smaller user counts. Tableau's products often fill an unmet need in organizations that already have a BI standard, and are frequently deployed as a complementary capability to an existing BI platform. Tableau is still less likely to be considered an enterprise BI standard than the products of most other vendors.

2.9 Open Source BI Software vs. Commercial BI Software

In the Open Source Business Intelligence context, we can see that most of the major vendors have free-of-cost community editions and enterprise versions, even some with free trials so you can try before you buy.

There is a community open source version with a well-defined set of functions, bounded and fully operational; and a professional version that presents more features or an enhanced version of the same features.

It may be that the free versions will meet your needs, if you have programming talent in-house able to customize your chosen BI software, but will need a big maintenance effort.

For instance, Pentaho Dashboards supports creating, but only the professional version has a Dashboard Designer Ad-hoc. These (premium) functionalities can be accessed only by purchasing a subscription or support.

2.9.1 What is open source business intelligence?

Open source BI are BI software can be distributed for free and permits users to modify the source code. Open source software is available in all BI tools, from data modelling to reporting to OLAP to ETL.

Because open source software is community driven, it relies on the community for improvement. As such, new feature sets typically come from community contribution rather than as a result of dedicated R&D efforts.

2.9.2 Advantages of open source BI tools

✓ Easy to get started:

With traditional BI software, the business model typically involves a hefty startup cost, and then there is an annual fee for support and maintenance that is calculated as a percentage of the initial purchase price. In this model, a company needs to spend a substantial amount of money before any benefit is realized. With the substantial cost also comes the need to go through a sales cycle, from the RFP process to evaluation to negotiation, and multiple teams within the organization typically get involved. These factors mean that it's not only costly to get started with traditional BI software, but the amount of time it takes is also long.

With open source BI, the beginning of the project typically involves a free download of the software. Given this, bureaucracy can be kept to a minimum and it is very easy and inexpensive to get started.

✓ Lower cost:

Because of its low startup cost and the typically lower ongoing maintenance/support cost, the cost for open source BI software is lower (sometimes much lower) than traditional BI software.

✓ Easy to customize:

By definition, open source software means that users can access and modify the source code directly. That means it is possible for developers to get under the hood of the open source BI tool and add their own features. In contrast, it is much more difficult to do this with traditional BI software because there is no way to access the source code.

2.9.3 Disadvantages of open source BI tools

- Features are not as robust:

Traditional BI software vendors put in a lot of money and resources into R&D, and the result is that the product has a rich feature set. Open source BI tools, on the other hand, rely on community support, and hence do not have as strong a feature set.

- Consulting help not as readily available:

Most of the traditional BI software - MicroStrategy, Business Objects, Cognos, Oracle and so on, have been around for a long time. As a result, there are a lot of people with experience with those tools, and finding consulting help to implement these solutions is usually not very difficult. Open source BI tools, on the other hand, are a fairly recent development, and there are relatively few people with implementation experience. So, it is more difficult to find consulting help if you go with open source BI.

2.10 Description of Barcelona and Genoa BI interfaces

2.10.1 Barcelona BI Interface

The municipality of Barcelona is now working with many BI softwares (Cognos, Microsoft, QlikView, Pentaho...) as a consequence of different strategies in different areas along the last period.

The last years have been of Cognos consolidation, although many departmental suites (as QlikView) have been maintained in order to supply specific needs with a fast-deployment solution.

Barcelona is now developing the future BI strategy at city level, analysing which are the best solutions to implement in order to choose the BI software that fits to the whole organisation with a broad vision, despite some departmental suites might still be maintained. It is expected to have a decision taken in the next 2-3 months.

The analysis being developed includes many comparative criteria that may be useful to share them with iCity partners.

- Gartner's Magic Quadrant: The vendors in the "Leaders" quadrant usually have the ability to combine good performance with a flexible response to market demands
- The most comprehensive tool for us: It is important to differentiate tools supporting all BI functionalities from specific products focused in an area (for example data analysis)
- Licencing costs - price to quality equilibrium: Analysis of functionalities needed and users to provide the solution to find the best balance for the organisation.
- Integration & scalability - previous platform integration: Look for the existing product that allows maximum integration with the existing platform providing scalability.

- Maintenance cost: Apart from the cost of licenses there is another cost to bear in mind as are the resources consumed by each one of the platforms (memory...)
- Market knowledge - Price negotiation advantages: The more knowledge has the market about a tool, the easier may be to find suppliers with expertise and best price competition
- Importing external data: Validate if the chosen tool has some complexity to import external data that may require technical background

2.10.2 Genoa BI Interface

The municipality of Genoa is now working with two different BI softwares:

- The first and most consolidated is based on the instruments supplied by Microsoft technology. This infrastructure has been used since 2006, and refers to Analysis Services for the multidimensional DB and Reporting Services for the presentation layer.
The choice of these products is due to the fact that Municipality of Genoa uses Microsoft sqlserver as a standard.
- Actually, to provide fast-deployment and distributed solutions we are implementing new environments based on Qlik View. This tool will become the standard BI solution provided to all the Municipality.

Both instruments will continue to be used by Municipality, relating to the different users needs.

2.11 DWH key points

2.11.1 Relational or non-relational Data Base

We need more than just the traditional relational SQL database functionality (RDBMS) for our iCity Platform.

NoSQL can service heavy read/write workloads compared to traditional RDBMS (relational database management system), and scale up to Terabyte (TB) size across replicated, commodity hardware. The primary driver for NoSQL is achieving massive, redundant distributed storage for large scale iCity applications.

If you want to make fast data analysis, you need to use a 'schema free' database architecture, designed for large data warehousing and high frequented data base changes.

A good example of a NoSQL database is the MongoDB. It can track and store tons of data about user interactions.

For example, you take an action on Meet Up, it will update your user references and update all your friends. **Non-relational stores** are really good at that and you can afford to keep that data in multiple places.

MongoDB (from "humongous") is an open source document-oriented database system developed and supported by 10gen. It is part of the NoSQL family of database systems. Instead of storing data in tables as is done in a "classical" relational database, MongoDB stores structured data as **JSON-like** documents with dynamic schemas (MongoDB calls the format **BSON**), making the integration of data in certain types of applications easier and faster.

For transactions, we still prefer a relational database, for example MySQL.

In a high scalable web architecture, a widely used open source memory object caching system is **memcached**. This is intended to speed up dynamic web applications by alleviating database load. It's a short term memory for the applications. (<http://memcached.org>)

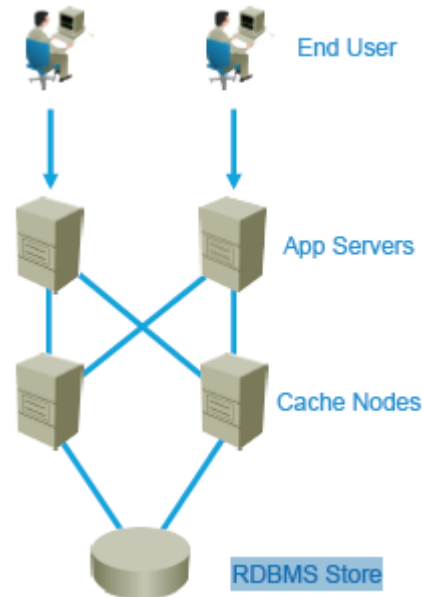


Figure 5 Dedicated memcached machines

An **Index store** is used as a cache for the meta-data: for example uploaded videos.

The index store reduces the need for complex and resource intensive queries to the RDBMS Store. (Multiple instances are required per deployment for redundancy)




Figure 6 Index Store

2.11.1 Video management

This section presents a proposal of video solution based on Cisco Show and Share³ product, that could be included in future prototype versions.

³ http://www.cisco.com/en/US/prod/collateral/video/ps9339/ps6681/data_sheet_c78-565776.html




Figure 7 Show and Share components

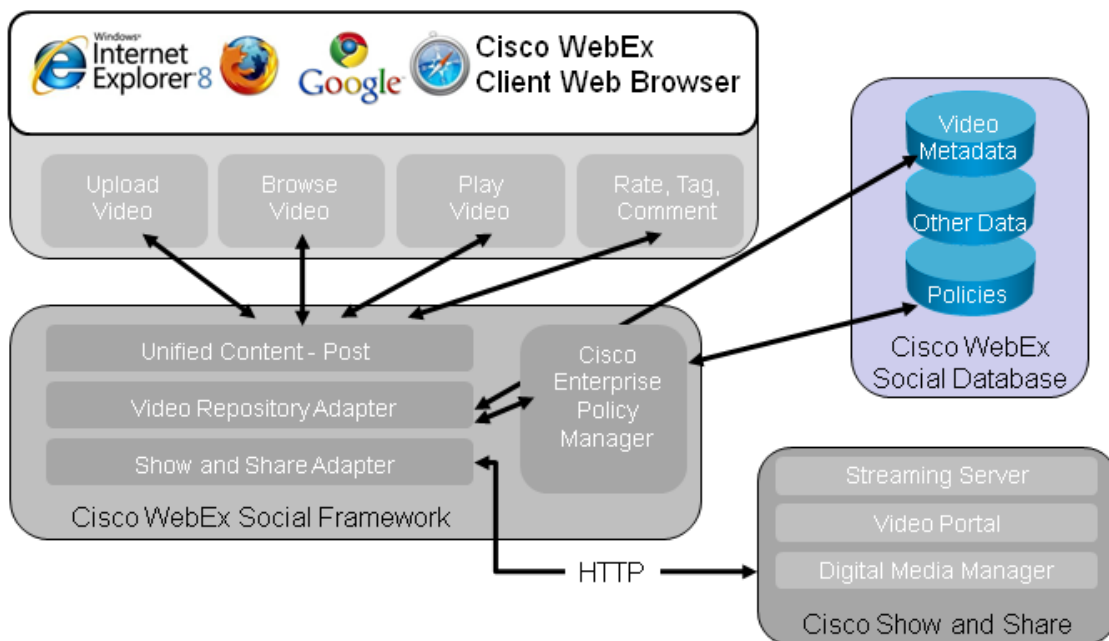
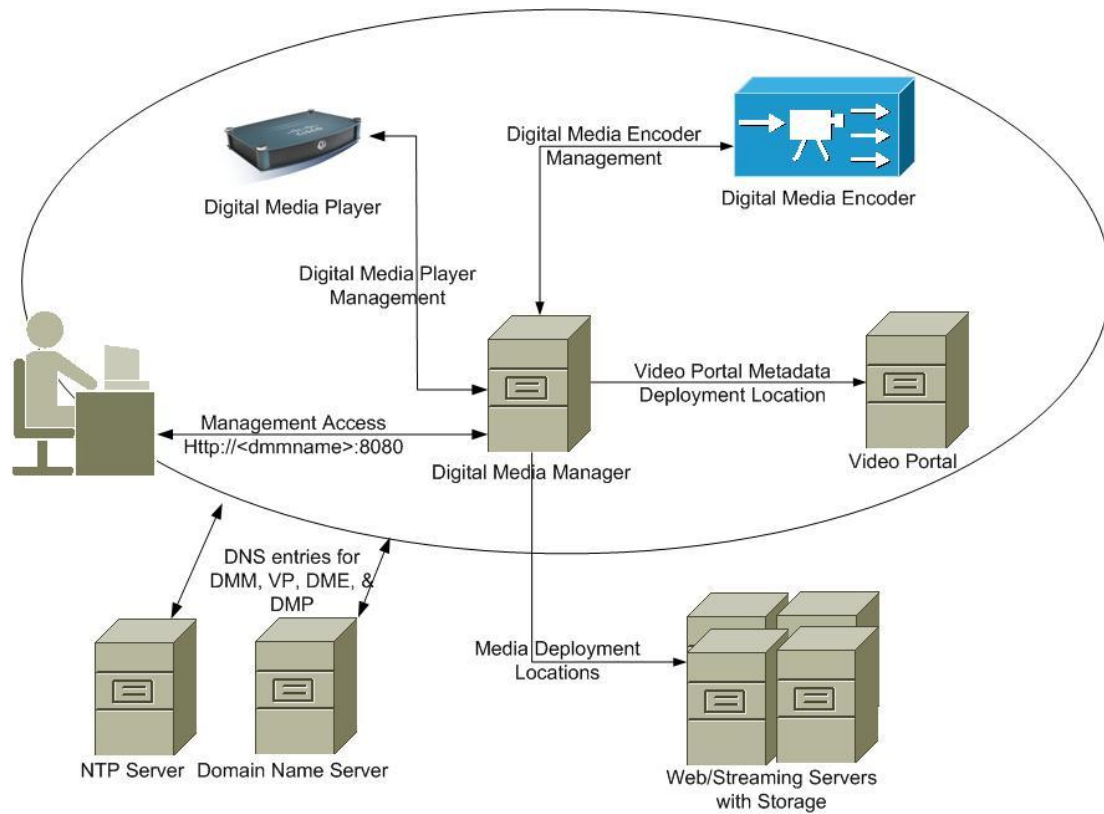


Figure 8 Show and Share architecture from client perspective



Used for videos embedded via the Webex Social Editor (e.g.: Posts, Message Boards). There will be APIs to integrate with other systems.

3. iCity DWH & BI Prototype

3.1 DWH adaptation

The iCity platform prototype adapts a simple version of data warehouse which overspreads the following requirements:

- ✓ Provide a standard protocol to access the data
- ✓ Resolve queries related historical data
- ✓ Be able to support future business intelligence deployments

The DWH adaptation prototype has been done to provide the following features, which will allow easy developments in the future:

- **Flexible** allowing the integration of different types of infrastructures located in the cities. The information provided by cities infrastructure is stored in homogeneous system that provides standard interactions for any service provided by iCity Platform. Finally, add new infrastructure could be done in an easy way.
- **Scalable** allowing to be improved and addition of new features in a easy and optimal way.
- **Dynamic** allowing the adaptation of different environment in order to resolve new requirements and needs.
- **Homogenized** allowing the iCity service to have a standard view and normalized access to data which is coming from resource layer.
- **Interoperability** providing technological transparency and allowing the communication of different technologies.
- **Modularity** based on quality software policies where different parts and procedures of the system are able to be reused, updated and replaced in a efficient and optimal way in order to facilitate the management.

Thus, first version of data warehouse is related with the storage layer and the catch up services module of iCity platform prototype as it is showed in the following picture:

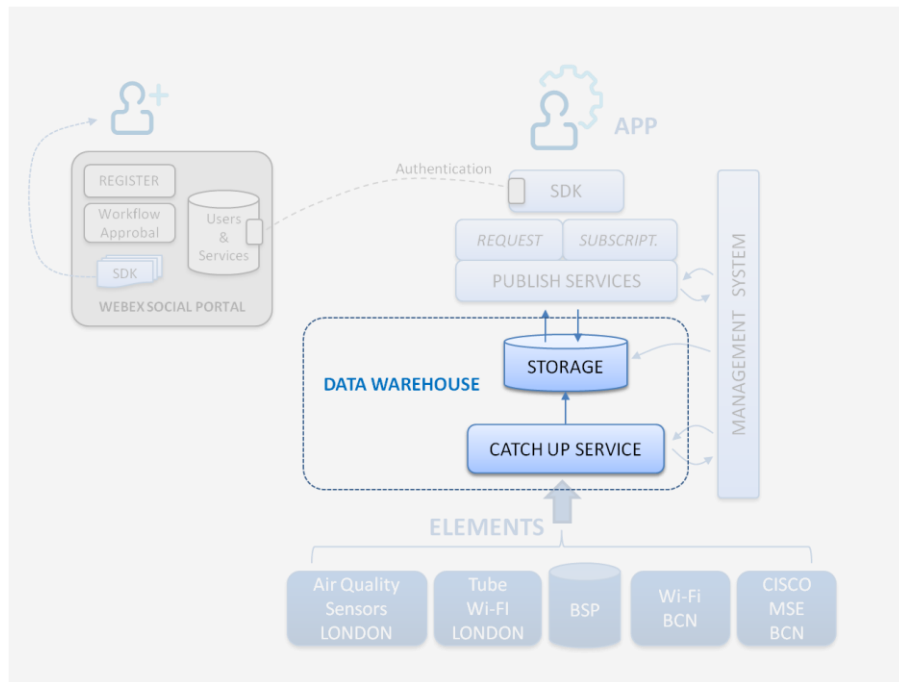


Figure 9: Data Warehouse prototype

Catch Up Services module:

This module is in charge of transferring the incoming data, which comes from the heterogeneous infrastructures, to the homogeneous storage layer of the iCity platform.

Cath up services acquire and recollect data from different information sources, allowing different communication protocols and offering a huge range of possibilities to integrate different types of infrastructure. Catch up service is composed of different modules, each module translates data coming from the infrastructure to a normalized model, reducing the integration and providing a generic communication of upper layers.

Finally, this module feeds the storage module with normalized data.

Storage module:

The storage module has the data structure based on OGC⁴, in order to provide homogeneous data to the upper layers of iCity platform prototype.

⁴ <http://www.opengeospatial.org/>

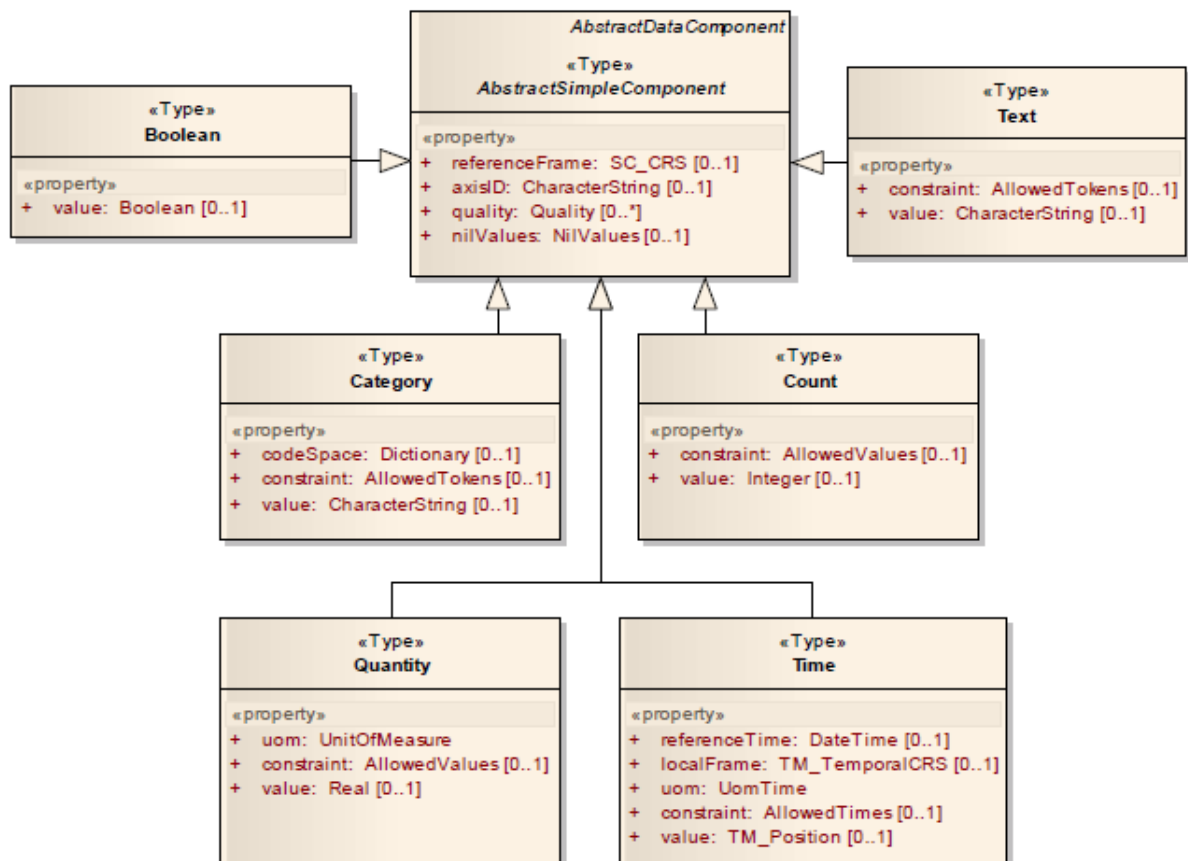


Figure 10: Example of data base model

Data coming from cath up services is managed by this module. Main features of this module are:

- To enable enrichment process that generates information from data, converting raw data to a cooked data.
- To provide normalized access to data for service layer in the iCity Platform.
- To offer a repository of historical data in order to reduce cost acquiring data from cities infrastructure.

At this stage of the project, as there isn't large amount of data to be managed, the storage module has been now developed using a relational data base.

4. Conclusions

This first year, a preliminary system has been included in the prototype in order to allow, in a short time, the deployment of pilots and also the development of apps, with a BI that would enrich and enable the validation of different exploitation business models and new business services.

Data Warehouse offers the following benefits:

- To integrate data form multiple sources.
- To perform new types of analyses.
- To reduce cost to access historical data.
- To normalize data across the iCity Platform.
- To share and allow others to easily access data.

Future evolutions of the prototype (D4.8 & D4.13) will follow inputs coming from WP3, WP5 and WP7.