BOARD Technology Overview

CONTENTS

1. PLATFORM OVERVIEW

1.1 BI, Analytics and Performance Management All-in-One

BOARD is the global leader in the all-in-one approach to Business Intelligence, Analytics and Performance Management.

From a functional standpoint, the all-in-one approach essentially means that BOARD provides in a single unified environment all the capabilities needed to build, manage and maintain any analytic solution for:

1. *Reporting and Interactive Dashboarding*
2. *OLAP Analysis and Data Discovery*
3. *Predictive Analytics and Simulation*
4. *Budgeting, Planning and Forecasting*
5. *Profitability Modeling and Financial Consolidation*
6. *Scorecarding and Strategy Management*

BOARD All-in-One approach

BUSINESS INTELLIGENCE

Scorecarding
Dashboard
Reporting
Analysis

PERFORMANCE MANAGEMENT

Budgeting, Planning &
Forecasting
Profitability Analysis
Strategy Management
Financial Consolidation

PREDICTIVE ANALYTICS

Forecasting
Simulation
Clustering
Statistical Functions

board

ERP, CRM LEGACY SYSTEM

RELATIONAL DATABASES, DATA WAREHOUSE

CLOUD, BIG DATA

EXCEL, TXT

OLAP CUBES

By unifying Business Intelligence, Analytics and Performance Management into a single product, BOARD merges data analysis and simulation with the planning, workflow and controlling cycles of an organization, enabling the transformation of data into insights, insights into simulations, simulations into plans and plans into traceable and monitored actions at a strategic, financial and operational level.

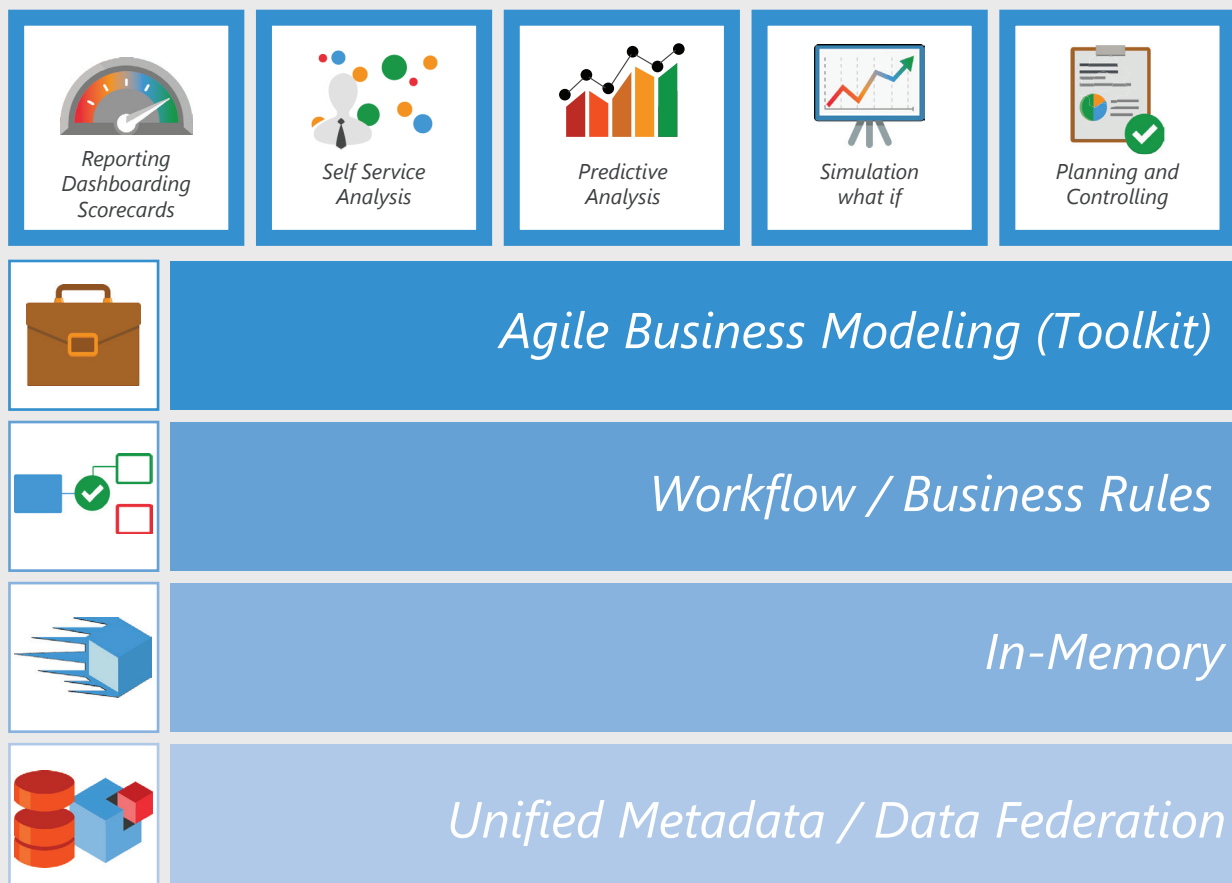This way, BOARD helps companies manage and control the entire decision-making process: from data collection to information analysis; from goal-setting to decision-making; from operational execution to monitoring results.

From a technical standpoint, the all-in-one approach requires the capability to run all these analytic, planning, simulation and monitoring activities into one unified platform built from the ground-up with:

1. *Single logical view of data (across BOARD MDB and external RDBM)*
2. *Single and shared metadata across any kind of solution delivered (BI, PM, Analytics)*
3. *Single User Interface*
4. *Single Security*
5. *Single administrative environment*
6. *Single Visual Modeling development environment*

Range of capabilities needed for delivering BOARD "All-in-One" approach

Reporting Dashboarding Scorecards	Self Service Analysis	Predictive Analysis	Simulation what if	Planning and Controlling

Agile Business Modeling (Toolkit)

Workflow / Business Rules

In-Memory

Unified Metadata / Data Federation

The unified data and metadata environment is the core foundation of the BOARD platform, but to fully explain BOARD's effectiveness in supporting decision-making processes it must be considered alongside three further essential capabilities which are required to complete the picture:

a) **A revolutionary and powerful Hybrid In-memory technology (HBMP)** *to manage information at the enterprise level and to empower navigation and analysis from a high-level financial and strategic perspective down to extremely granular operational data.*

b) **An agile business modeling environment (The Toolkit)**, *to easily build customized solutions that perfectly fit the organization's requirements and to instantly adapt them to ever changing business needs.*

c) **A fully integrated workflow and business rule back-end,** *to build digitalized analytic, planning, monitoring and simulation processes, supporting collaborative decision making across the entire enterprise.*

TOOLKIT

Drag and Drop
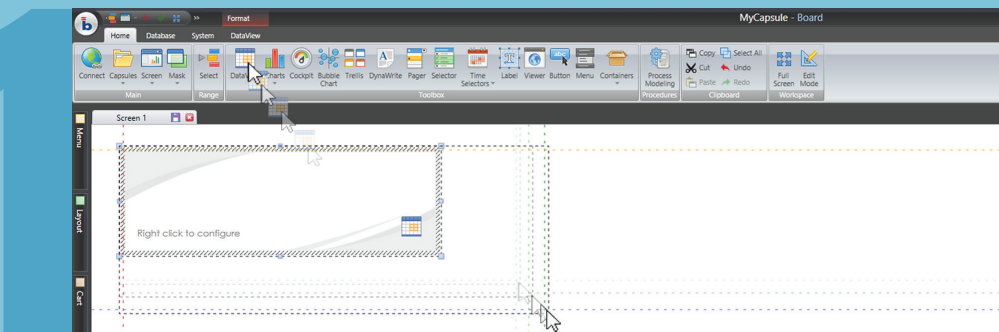
Configure
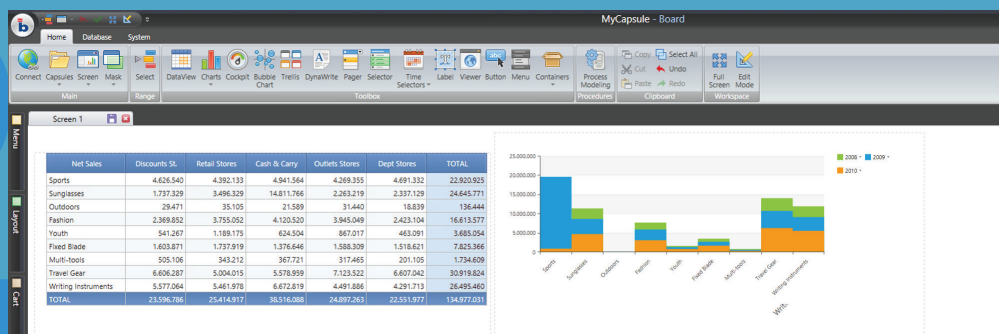
Run

1.2 The Toolkit: the fastest way to deliver analytic solutions

BOARD provides the user with every tool needed for database creation and update, data presentation and analysis and process modeling, in a single visual and interactive environment, making it easy to quickly build customized analytic applications ("Capsules").

A "Capsule" consists of a series of screens that can accommodate any BOARD object: reports, graphs, spreadsheets, folders, point and click filters, simulation models, gauges and navigation menus.
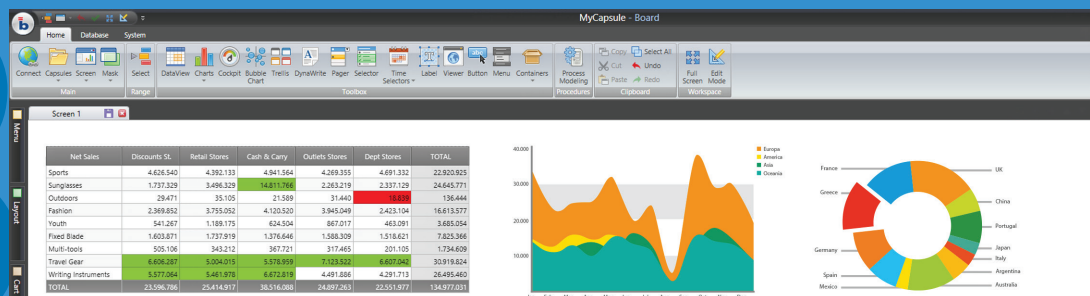
These objects can be freely positioned on the screen by simply dragging and dropping from the ribbon bar, then interactively configuring using the comprehensive settings and format options.

They are automatically synchronized with each other and with the database, allowing users to instantaneously deliver and visualize any development. Furthermore, a powerful set of business rules allow the user to work on objects, data and workflow to model all the typical business processes (multi-dimensional data entry, calculation, allocation, consolidation, etc.) needed to create performance management applications.

1) Eliminates the need for high levels of technical expertise

2) Enables perfect customization of any application to meet the business need

3) Provides extreme flexibility in responding to ever changing business needs

4) Perfectly fits the "Agile Development" methodology

5) Substantially reduces the complexity and costs associated with building, customizing and maintaining a comprehensive BI and Performance Management environment.

1.3   Self-service analysis: insights at users' fingertips

BOARD enables end users to create personalized reports and analytical queries, freeing up IT staff to focus on tasks that are more critical.

An innovative user environment that combines search based data discovery, personal dashboard mash-ups and drag and drop functionality makes access to the right information easier, faster and more effective for everyone.

The integrated search empowers users to easily explore data, metadata and applications and to use results as drivers for further exploration. Once the right information is identified, the drag and drop environment enables anyone to immediately attain personalized analyses and reports. Finally, the pinBOARD, enables users to save their findings and to create their own personal analytical environments.

These powerful self service analytics do not simply change what the users can do, but drastically transform and enhance how they can do it.

Furthermore, thanks to the native Business Intelligence and Performance Management unification, this new capability extends to every type of analytic application built with BOARD, closing the gap between the individual need for self-service analysis and the organizational necessity to have a shared vision of the truth throughout the whole enterprise.

BOARD Self-Service analytics

Search *information*

1 **BOARD** *Integrated Search*

Customize *& Analyze*

2 **BOARD** *Drag and Drop Analysis*

Make it **yours**

3 **BOARD** *PinBOARD*

1 BOARD Integrated Search

BOARD offers an interactive search function working across screens, data, metadata, and applications. As a single, unified solution, BOARD allows any search result to be immediately used as a driver for further analysis (e.g. if the search is for a product name the user can immediately choose a report where this is included and use the product or a co-related dimension as navigation criteria) or as an element for creating new reports, using the new drag and drop self-service analysis user interface.

Search
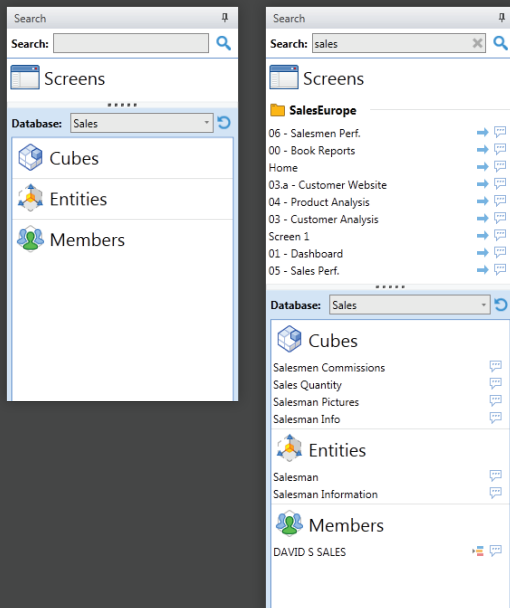
Navigate search results

1 Screens
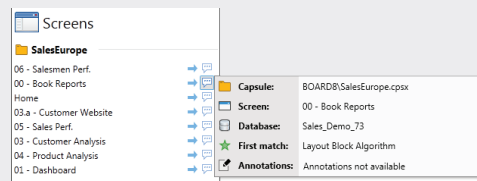
2 Cubes
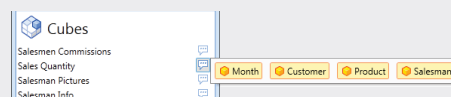
3 Entities
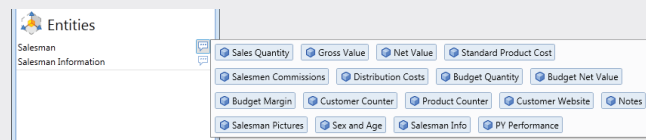
4 Members
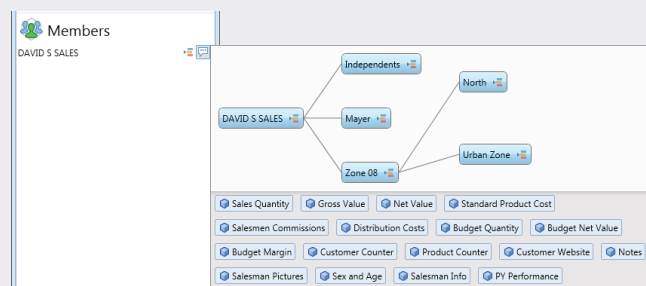
2 *Drag and Drop analysis*

BOARD gives users not only the ability to create their own reports on the fly, but also to personalize centralized reports using an intuitive drag and drop interface. This way, any user can choose values, set rows and columns, define calculations, run reports and get business insights at the speed of thought.
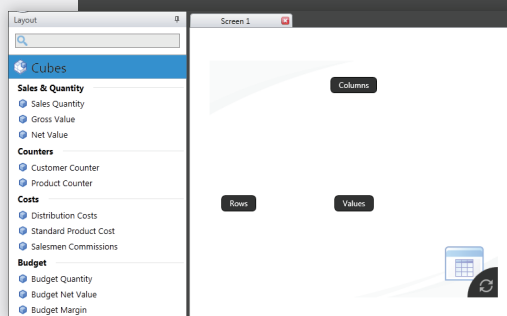
1 Access the drag and drop sliding panel

2 Select values
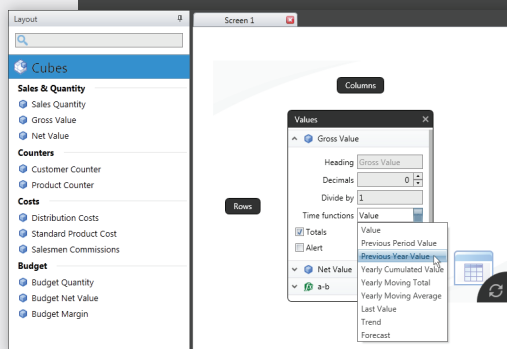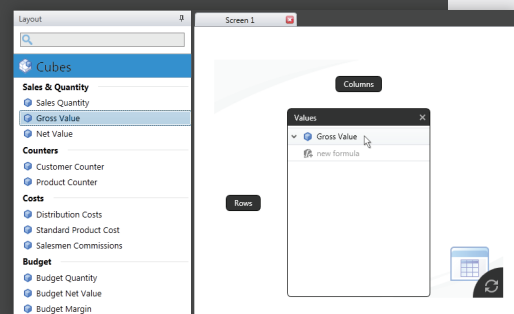
3 Define formulas and time functions

4 Set rows and columns

5 Run the report

3 The PinBOARD: your personal analysis environment

BOARD enables each user to build their own personal Business Intelligence applications. Users can select and combine any existing report, graph or gauge and save it as an application into a personal environment called "PinBOARD". All the reports can also be quickly personalized through the drag and drop functionality, all without the need for IT to intervene. This innovative mash-up environment is an easy and smart way to offer any user personalized access to information that is physically distributed across the entire enterprise, without compromising data security, reliability and integrity.

🛒 Add to cart

📌 PinBOARD

1.4　*Cutting edge technology*

Since its foundation in 1994, BOARD has pioneered a unified approach to Business Intelligence and Performance Management which we have called "Management Intelligence". This history makes BOARD a unique solution in the BI and Performance Management market. In fact, unlike its main competitors, BOARD is not a patchwork of different products and technologies coming from multiple acquisitions and forcibly integrated, but is instead the result of an organic development project nourished over years by the continuous flow of feedback coming from over 3,000 customers.

Furthermore, beginning in 2007, BOARD started a massive software development effort to rewrite the whole product, utilizing the most advanced technology available on the market. In 2009 a brand-new version of the product totally based on the Microsoft .NET framework and WCF was released. Currently BOARD is aligned with the latest version (4.5) of this environment which is rapidly becoming the leading platform for building and running next generation applications.

From a functional standpoint, the adoption of this new technological platform permitted a drastic improvement in the way in which BOARD was engineered and developed and, at the same time, created the conditions necessary to develop new functionality enhancements and innovations at a fast pace.

From an architectural standpoint, the adoption of the Windows Communication Foundation made it possible to start offering an innovative SOA architecture with a distinct separation between business and data presentation logic and enabled a seamless bi-directional integration with the Office environment.

From a business standpoint, the introduction of modern and widely recognized standards guarantees to any BOARD customer that they are investing in a future-proven and long-term oriented technology, safeguarding the value of their money.

The underlying technology, though important, is not the most significant part of the picture. In fact, BOARD LABS have been pushing the boundaries of innovation in the BI and Performance Management space for years.

From the creation of the most agile multi dimensional database on the market to the introduction of the Toolkit (Visual Application Development); from the Reverse Algorithm (cells that support formulas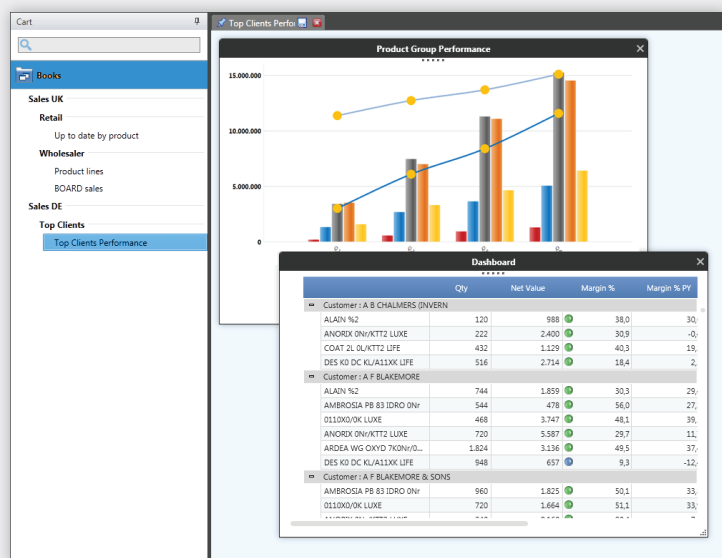 and data-entry at the same time) to the ATO (real time data modeling); from the HBMP (Hybrid In-Memory with write-back) to the PinBOARD (personal analysis environment), countless unique features and technologies have always characterized BOARD as one of the most innovative and technologically advanced platforms available on the market.

The start of a complete **redesign** of the product based on brand new technology

1994	1996	1997	2000	2005	2006	2009	2010	2011	2012	2013	2014
RELEASE OF BOARD 1.0	**REVERSE ALGORITHM**	**INTRODUCTION OF THE TOOLKIT**	**FIRST BOARD WEB VERSION**	**ROLAP** TECHNOLOGY	**ATO OBJECT**	RELEASE OF **BOARD 7.0**	**ADVANCED DATA ENTRY**	**HBMP TECHNOLOGY**	**SELF-SERVICE ANALYSIS**	**MOBILE DATA ENTRY**	**PREDICTIVE ANALYTICS**
The first software platform to combine Business Intelligence with Planning	Capability of data entry on calculated field (like if an Excel cell could support at the same time a formula and a value, making it possible to use it in output and input)	Release of **BOARD 4.0** Visual development of any BI and Performance Management application	Release of **BOARD 5.0** Introduction of BOARD HTML Client	Release of **BOARD 5.5** Capability to directly read and write back on relational database	Release of **BOARD 6.0** Dimensions members such as products, business units, initiatives etc. can be added to the data model during the planning process directly by the users, and immediately utilized as planning units.	**A brand new product version,** completely rewritten using the latest .NET technology	Release of **BOARD 7.1** (Cell freezing and advanced allocation of data changes)	Release of **BOARD 7.4** Introduction of a new database based on the **new frontier in in-memory BI & CPM**	Release of **BOARD 8.0** Introduction of data-discovery (Integrated search, drag&drop analysis, pinboard)	Release of **BOARD 8.1** Introduction of Native Mobile for iPad and W8 supporting data entry	Release of **BOARD 9** Seamless integration of advanced and predictive analytics into the BOARD platform

2. PRODUCT ARCHITECTURE

2.1   BOARD: functional architecture

BOARD is an enterprise-class platform based on a single unified n-tier architecture that offers:

1) *Scalability*
Any BOARD element from the multidimensional database to the application architecture has been designed and developed to ensure full scalability in terms of Number of Users, Data Volumes and Application Scale.
An integrated in-memory clustering technology makes it possible to support large (thousands of users) worldwide projects, avoiding geo-latency problems.

2) *Performance*
BOARD offers outstanding performance not just in terms of data visualization, but also in handling large and complex planning processes with thousands of data-entry concurrent users.
This is all made possible by the combination of a groundbreaking proprietary Hybrid in Memory technology (HBMP) with a powerful MDB that has always been a speed benchmark in the BI space thanks to its bottom-up structure and sophisticated sparse-data management technology.

3) *Consistent experience*
Any application built in BOARD can be accessed by multiple clients (Browser, BOARD client, Tablet, Office add-in) providing a consistent user experience to the business users.

4) *Robust security*
Security is pervasive across the whole platform architecture. From authentication to cell level authorization, BOARD offers a robust, comprehensive and easy to manage enterprise-class security system.

5) *24 x 7 availability*
Thanks to its data model flexibility (that allows users to make changes in real time) and to the ability to upload data while analysis or planning processes are running, BOARD makes it possible to deliver 24 x 7 Business Intelligence and Performance Management applications.

Invented, designed and developed as a single unified platform since its inception, BOARD's structure is composed of two main parts: the BOARD server and the BOARD Clients.

2.2   BOARD Server

The BOARD Server is the engine that carries out all processing and is based on three building blocks.

a. The Multidimensional Databases
Ensure the capability to access, normalize, federate, manage and model data.

b. The Business Logic Environment
Offers a comprehensive set of functionalities and business rules to implement, run and manage planning, budgeting, forecasting, consolidation, allocation and simulation logic – seamlessly linking data with business processes.

c. The Applications Environment
The BOARD end-user environment that allows:
-business users to view, interact and analyze data
-developers to quickly build and run any application without coding, leveraging all the functionalities/capabilities offered by the MDB and the Business Logic environment.

The BOARD server executes all aggregations, calculations, selections, procedures, data import processes and any other interaction involving the BOARD multidimensional database. It also handles incoming connections from BOARD users.
The BOARD server performs user authentication, applies security restrictions or privileges and then dispatches user requests as different execution threads to carry out the multidimensional processing.
Communication between the client and the server uses a proprietary protocol named ROAR (Remote Object Access & Replication) which provides extremely high performance in communication and is designed to work efficiently over low band-width connections such as WAN or the Internet.
The ROAR protocol is built upon the Windows Communication Foundation (WCF) classes of Microsoft .Net Framework 4.5.

2.3   BOARD Clients

BOARD ensures a consistent experience across web, desktop and mobile applications so that business users can rely on the same familiar environment on whichever device they are using.

The **BOARD Clients** are the diverse BOARD user interfaces that allow users to access BOARD applications (capsules), namely:

a. BOARD Desktop client

b. BOARD Web client

c. BOARD Mobile (iPad, W8 native apps)

d. BOARD MS Office Add-ins (Excel, Word, PowerPoint)

BOARD: Logical Architecture Schema

DATA SOURCES	BOARD ENGINE			CLIENTS
	BOARD DATABASES	BUSINESS LOGIC	APPLICATIONS	
RELATIONAL **DATABASES**, DATA **WAREHOUSE**	UNIFIED METADATA	DATAENTRY	CAPSULE **A**	CLIENT **WPF**
ERP, CRM **LEGACY SYSTEM**	DIMENSIONS	DATAFLOW		CLIENT **WEB**
CLOUD, BIG DATA	HIERARCHIES	ALLOCATION CONSOLIDATION, SIMULATION	CAPSULE **B**	OFFICE **ADD-IN**
EXCEL, TXT	CUBES	WORKFLOW		TABLET **MOBILE**
OLAP **CUBES**	HETEROGENEOUS DATA	ALERTING	CAPSULE **C**	PORTAL **MASHUP**
		BROADCASTING		

2.3.1   BOARD Desktop Client

Conceived for offering a rich Windows experience and all the power of a thick-client, BOARD Desktop is the most comprehensive BOARD user interface.

BOARD Desktop allows users to:

a. Access BOARD applications

b. Run interactive analyses

c. Execute data entry

d. Build application

e. System administration and management

A modern auto-update functionality makes it easy to use BOARD Desktop clients in large-scale deployments, eliminating all the technical issues usually associated with windows-client versioning.

2.3.2   BOARD Web Client

The BOARD Web client is a zero-footprint, rich-internet client based on Microsoft Silverlight. Microsoft Silverlight is a cross-browser, cross-platform technology that provides the BOARD web client with the ability to run in all popular Web browsers, including Microsoft Internet Explorer, Mozilla Firefox, Apple Safari, Google Chrome and on Microsoft Windows and Apple Mac OS X.
BOARD Web client allows users to:

a. Access BOARD applications

b. Edit reports and run interactive analyses

c. Execute data entry

From a technical standpoint, the BOARD Web client is a runtime environment which can run BOARD applications from any browser. This means that, once developed, any BOARD application can be accessed by either the Web client or the desktop client making it possible to manage hybrid (Web and client) deployments without having to build and maintain two different environments.
Furthermore, fully supporting mash-up, the BOARD Web client can make BOARD inputs

available to external applications or acquire inputs from external applications without the need to build integration adaptors.

This important feature, combined with the Single Sign On, enables BOARD to be easily integrated with any company portal.

2.3.3 BOARD Mobile

Designed and developed for a high performance, gesture based interaction, BOARD Mobile is a native environment for Windows 8 and iOS based tablets.

BOARD mobile allows users to:

> *a. Access BOARD applications*
>
> *b. Run interactive analyses*
>
> *c. Execute data entry*

Users can also access their data in offline mode making it possible to work without connection to a network.

Built using the toolkit approach, exactly the same way as standard BOARD capsules, BOARD Mobile applications allow organizations to deploy almost any type of screen to a mobile device in a matter of minutes.

Consistent User Experience

BOARD ensures a consistent experience across Web, desktop and mobile applications so that business users can rely on the same familiar environment on whichever device they are using.

Mobile Data Entry

BOARD Mobile client is the only BI and Performance Management solution to fully support data-entry on the go.

Users can insert data on their tablet, save and immediately have the changes reflected into their mobile reports and analysis. This allows organizations to realize next generation mobile business processes for simulation, planning and forecasting.

Shared Insights

BOARD Mobile makes it possible to share dynamic reports and analysis between BOARD Mobile users by sending links. Based on the security profile, any link receiver will be able to not only access shared reports and analyses, but to also interact and modify them as they like.

Furthermore, any analysis or dashboard can be easily shared as a static report on social networks through the native iPad/W8 functionalities.

Offline Access to Key Data
BOARD Mobile allows users to work both online and offline - users can select which reports and dashboards they wish to make available for offline use.

Rapid Development and Deployment
Using the programming-free BOARD Toolkit, mobile applications can be developed and deployed without a single line of code and without having to learn any new development environment or methodology.
All the applications built using the BOARD Toolkit can be effortlessly converted in mobile applications due to an automatic compatibility check for layouts and features that reduces development errors, time to solution and deployment issues.

2.3.4   BOARD MS Office add-ins

BOARD empowers users with self-service analysis and reporting in a familiar MS Office environment.
Through the BOARD MS Office add-ins, users can access BOARD multi-dimensional data directly within Excel, Word and PowerPoint and utilize the traditional OLAP analysis functions, such as drill-down, slice and dice, filtering and ad hoc query.
BOARD objects such as reports, charts and gauges can be easily embedded in Microsoft Office documents making it easy to create auto-updating booklets, factsheets and presentations.
The data flow between BOARD databases and Office documents is streamlined by an auto-update functionality that allows users to schedule data refresh at their convenience ensuring full automation of the entire process and elimination of any manual activity.
The BOARD Excel add-in also supports data-entry in online and offline mode offering an extremely effective option for managing data collection processes: users can work without being connected to BOARD and automatically synchronize data when they first reconnect.
All of this is delivered with all the security robustness and the data consistency of the BOARD platform.

3. SECURITY

3.1　Multi-tiers security model

BOARD's security model is made of multiple functional layers that can be thought of as independent tiers providing different abstraction levels where security policies, methods and definitions are maintained or integrated with other network or legacy systems.

The two main concepts are Authentication and Authorization. They are distinct layers, linked through the role-based security concept.

Authentication

1

USER'S AUTHENTICATION
• BOARD
• Windows Authentication
• LDAP
• SAML (PKI)

Authorization

2

ROLE-BASED PROFILE
• Consumer/Builder/Admin
• Read Only/Data Entry
• Export/Print/Mobile

3

APPLICATION SECURITY
• Capsules Access
• Capsules Editing

4

DATABASE SECURITY
• Cubes (Measures)
• Members
• Single Cell
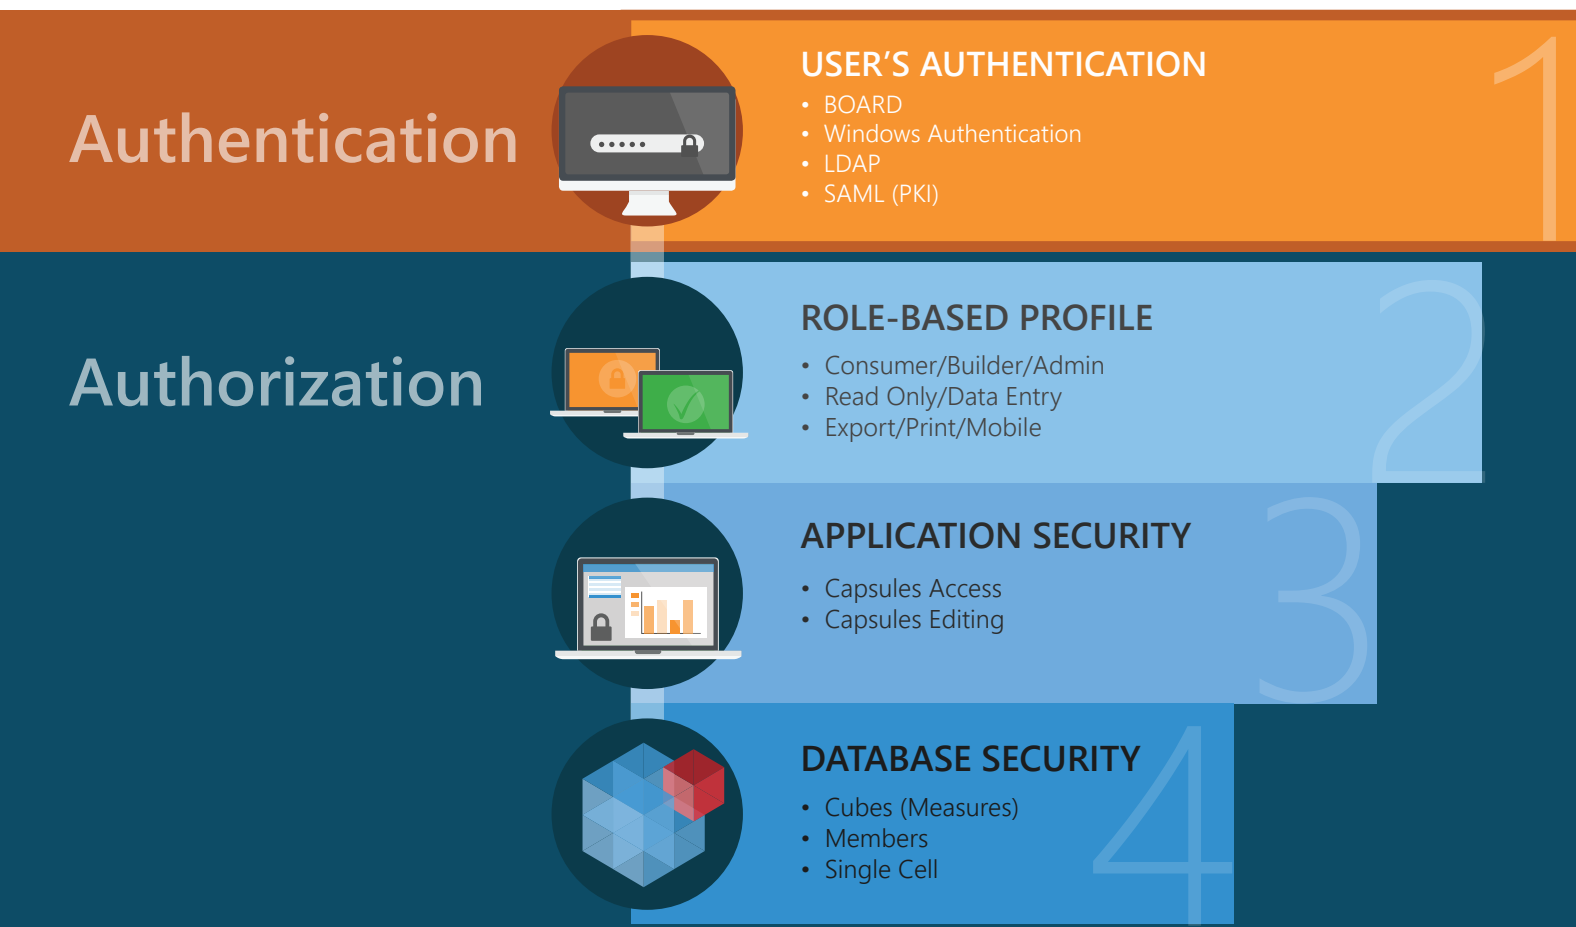
While Authentication is only performing username/password verification, the Authorization module is made of three layers corresponding to three logical levels at which authorization can be defined, explained hereafter.

3.2   Authentication methods

The Authentication layer supports multiple authentication methods which can be selectively enabled or disabled:

1) *BOARD authentication: this is the native built-in authentication. An administrator can create new BOARD accounts (i.e. username/password) and define rules such as password expiration, minimum length, maximum retries.*

2) *LDAP authentication: the username and password given by the user is validated against an LDAP server. This method only requires to configuring a connection between the BOARD Server and the LDAP server, for example LDAP://myserver.mydomain.com:389/un=username, dc=mydomain, dc=com This method also supports cross-domain authentication.*

3) *Windows Authentication: in a Microsoft Windows network, the BOARD Server can use the domain controller to validate the user identity. When this method is used, the BOARD Client, Office add-in client or Web client can reuse the current Windows session authentication providing a Single Sign-on with the active Windows user..*

4) *PKI authentication (SAML): with this authentication method, the BOARD Client uses SAML protocol to authenticate the user on an Identity Provider server and part of a Public Key Infrastructure (PKI). A PKI authentication typically provides smart-cards or one-time password (OTP) authentication methods.*

BOARD Authentication methods
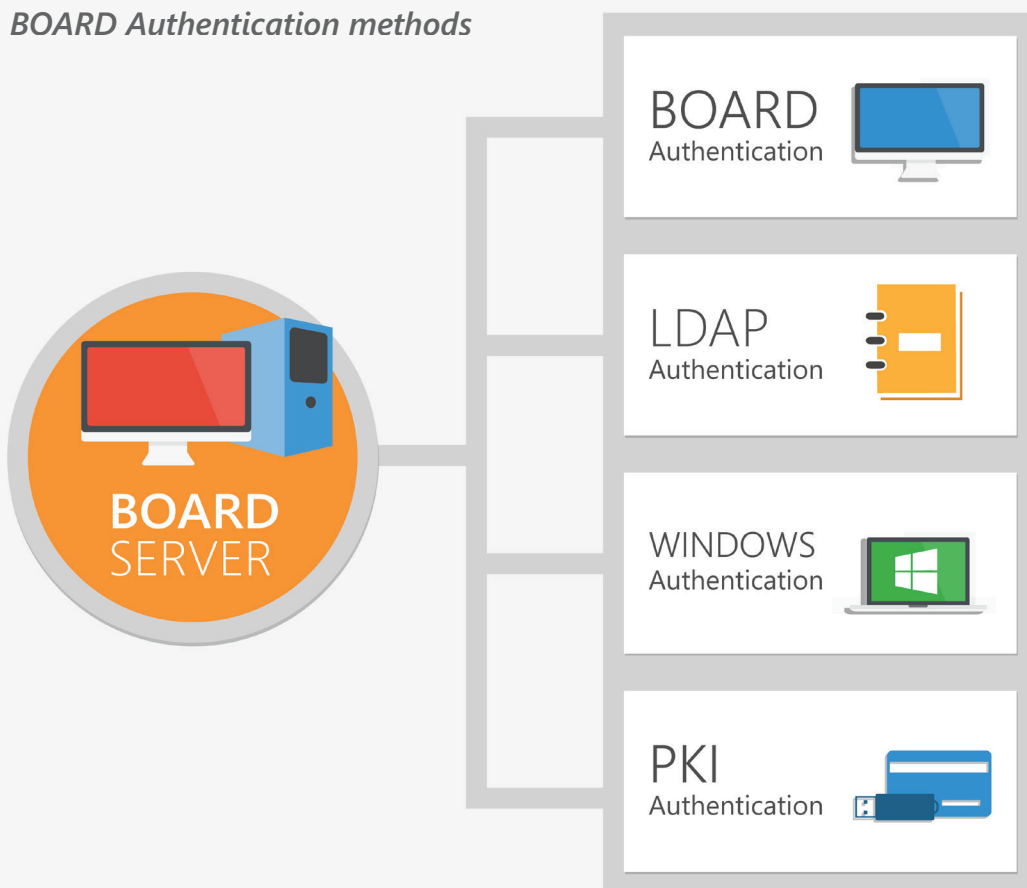
3.3   Role based authorization

A set of authorizations is stored as a BOARD security profile which can be assigned to individual user accounts. Following the role-based security paradigm, definitions of authorizations are not defined at individual user level but as a role. There are two main sets of authorizations which form a security profile: one defines authorizations on general features of the environment such as the Power User or Developer features and the other defines authorizations on data such as cubes and filters on dimensions.

In a BOARD security profile it is possible to define authorization to:

> *a. **Database and Capsule authoring features through the license profile***
>
> *b. **Restrict or grant access to given Capsule folders***
>
> *c. **Restrict or grant access to a BOARD database and with what database profile***
>
> *d. **Restrict or grant access to filtering, printing and other features***

3.4   Application authorization

The BOARD Capsules can also embed additional authorizations (for example to allow opening or editing of a Capsule). This level of authorization ensures that a Capsule can't be opened or altered even if it is copied or moved to a different environment (i.e. a different BOARD Server) where the user may have a higher set of privileges.

The application level authorizations can also be used to restrict execution of critical actions such as running a certain procedure that processes data in the solution.

3.5   Data access

The authorizations defined on the data model are stored in the BOARD Database and allow:

> *1) Granting/denying access to create or modify a data model: authoring dimensions, cubes, data sources and other database objects,*
>
> *2) Granting/denying access to cubes with read-only or read-write privileges*
>
> *3) Setting filters on dimensions down to cell-level (for example to limit a certain security profile to a subset of data using the entities of the data model, such as a Region or Company, a set of Cost Centers or a Division)*

3.6   Advanced security management

BOARD offers the possibility to store and manage authorizations in a BOARD database that contains the metadata of all other BOARD databases. A simple idea - use BOARD to manage information (metadata) about BOARD databases and solutions and therefore build a completely bespoke solution to handle security grants using BOARD tools and features.

This capability not only makes user management easy and efficient, it also allows the delegation of security administration to key users, making it possible to decentralize the security model and to perfectly adapt it to complex organizational structures.

BOARD Advanced security management
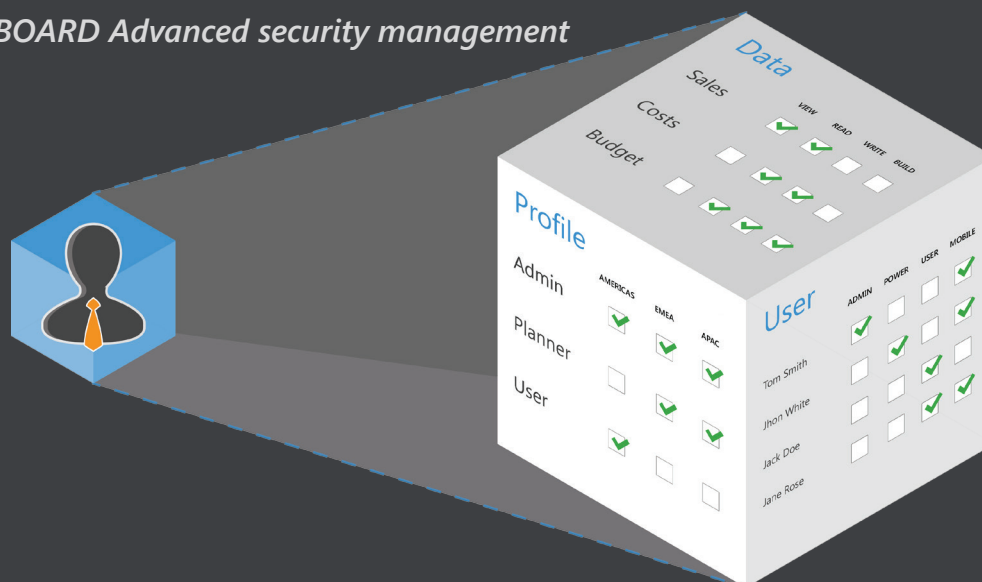
This enterprise feature uses a special BOARD database where cubes are used to grant authorizations and access rights of the Security Profiles or of Users. This permits Segregation of Duties (SoD) where more than one person is required to authorize an account or a given authorization level.
Creation, deletion, modification of authorizations or authorization levels are logged for auditing purposes.

4. DATA INTEGRATION

While in theory the aim of any company is to maintain data stored in a single enterprise data-warehouse, often business complexity makes this too demanding, difficult or even impossible. In everyday business life, large volumes of data are stored in a wide variety of sources and platforms throughout the enterprise, and even outside (cloud).

The BOARD platform provides the capability to connect, integrate and federate data across:

a. Relational databases and Data Warehouse(s)

b. Enterprise applications (e.g. SAP ERP)

c. Multidimensional sources (including SAP BW)

d. Web Services

d. Excel, CSV and txt files

d. Cloud based sources (e.g. salesforce.com, OData)

Through its physical and logical Multidimensional Databases, BOARD offers a single metadata gateway that makes it possible to use data as if they were native to BOARD.
Data is normalized and organized in BOARD cubes, allowing end-users to read, write and update them regardless of the data sources.
Behind the scene, the administrators have the freedom to decide which data to store into BOARD multi-dimensional databases (MOLAP), and which to leave in place and manage through BOARD data-federation capabilities (ROLAP).
The BOARD DUAL Engine will automatically manage the two different options, providing business users with a single and seamless vision of the data.
Furthermore, unlike its main Business Intelligence competitors, BOARD supports write back not only on its cubes, but also directly on relational data sources making it possible to effectively integrate performance management processes with enterprise applications (i.e. storing forecasting/planning data directly inside the ERP).

From a technical standpoint BOARD uses:

1) Open Database Connectivity (ODBC) standard and OLE DB to connect with relational databases

2) OLE DB for OLAP (ODBO) to connect with multidimensional database.

3) ODATA to connect to Cloud Based sources

4) CSV, TXT files

4.1   Multidimensional data sources

BOARD supports OLE DB for OLAP (ODBO) and XML for Analysis, providing customers with a standard method of accessing the most popular multidimensional engines, such as:

a. Microsoft Analysis Service

b. Oracle/ Hyperion Essbase

c. SAP BW

d. IBM Cognos TM1

ODBO (OLE DB for OLAP) is a Microsoft specification and an industry standard for multi-dimensional data processing specifically designed to access multi-dimensional data sources. XMLA is essentially a newer standard interface. Both of these technologies operate using the Multi-Dimensional Expressions (MDX) language to query cubes.

BOARD Data Integration
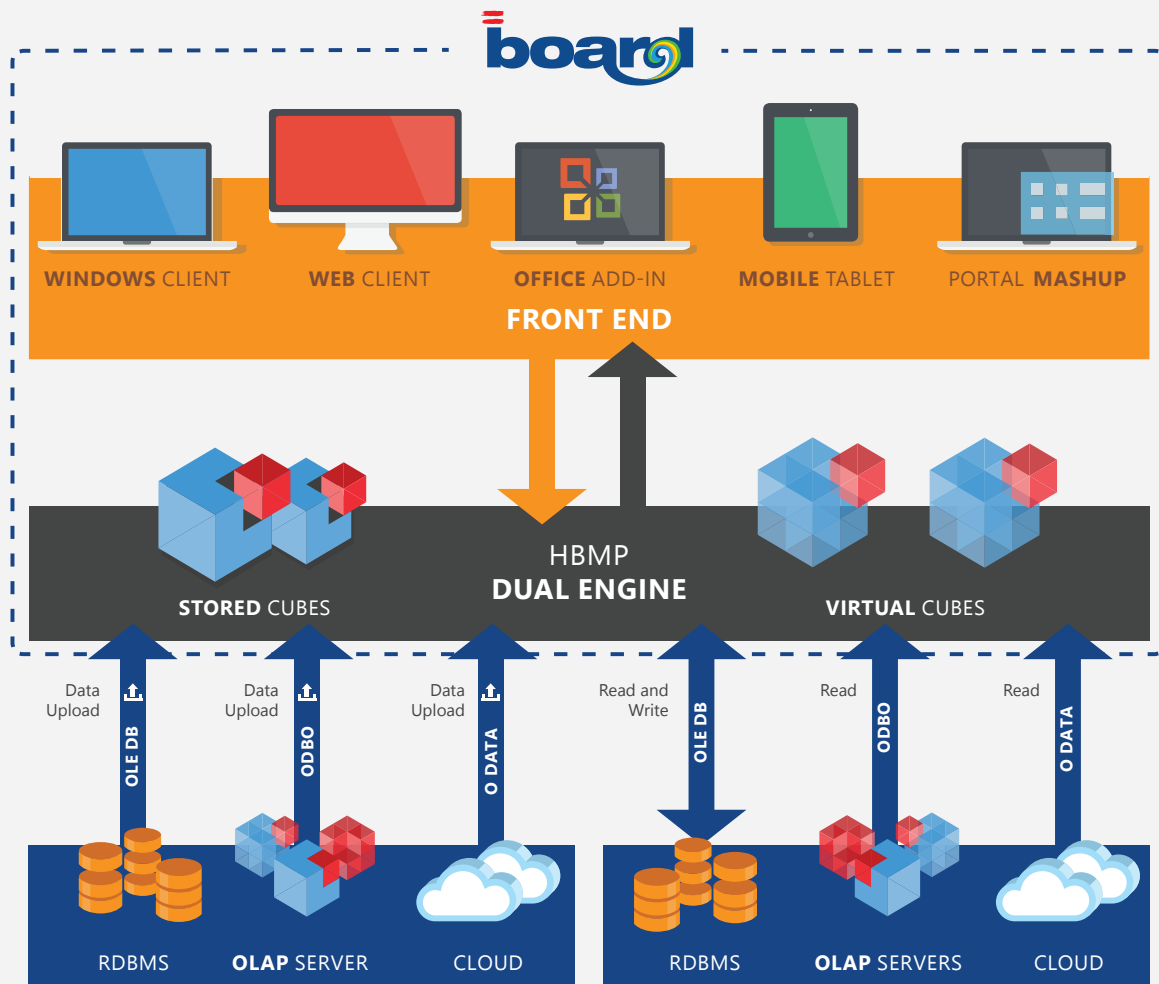
4.2   Data Federation

Data Federation is the process to gather data from distinct databases without transferring the original data itself. Rather than collect all the information in a database, physical data is left in place, and metadata - data that describes the data itself and its location - are utilized to create a virtual database.
Through its Data Federation functionalities (ROLAP Engine), BOARD offers the capability to read and write data stored in distinct RDBMS and to use them in a transparent way, as if they were stored in the BOARD Multi-dimensional database.

More precisely, metadata is stored in BOARD's virtual cubes named "RDB cubes": whenever a user views a report based on these measures, an ODBC or an OLE DB connection to the external relational database is established and data is retrieved from the relational table through a SQL statement.

This way anytime a user runs a multisource request, BOARD automatically:

a. Generates a series of queries for each database needed by the report/analysis

b. Connects to the corresponding data sources for execution

c. Feeds the BOARD virtual cubes with the final results from each of the queries

d. Transparently makes the final merged results available in the report/analysis

From an architectural standpoint, this capability allow developers:

1) To build extremely efficient BI and CPM applications by eliminating the need to upload data at the lowest level of detail into BOARD Databases without compromising the capability to explore them.
In fact, data can be physically stored at the most suitable aggregation level for each analytic model thus reducing data volumes and time-consuming upload processes while granting the capability to access the lowest level of detail through virtual cubes whenever needed.

2) To create near-real time analytic models by making RDBMS data directly available in the BOARD MOLAP environment.

4.3   Essential ETL functionalities

Despite not being a dedicated ETL product, BOARD offers capabilities that make it possible to directly feed BOARD databases from diverse source systems, such as an ERP a CRM or other legacy systems, without the need for intermediate data staging layers.

This is unique in comparison to most other Business Intelligence products which typically require the source data to be cleansed and organized into either a star or snowflake schema.

This can be a significant cost during the implementation and a substantial limit to the flexibility of the analytical applications that often require working on data source organization even for simple changes to the data model.

The key ETL capabilities provided by BOARD through its ETL and Data Reader component are:

1) *validation rules*
defining simple or composite validation formulas/criteria such as:
a. *ignore those records where a certain key or value is not present*
b. *ignore or accept records where certain fields contain a given value (or threshold or set of values)*

2) *normalization and conversion*
a. *translating coded values, e.g. source system contains 1 of Male and 2 for Female but we want to use M for men and F for women*
b. *encoding free-form values, e.g. mapping DE to Germany, FR to France*
c. *standardising codes from multiple sources, e.g. one source system uses 01 for a given product and another source system ABC then use one single unique code.*
d. *deriving a new calculated value, e.g. sale_amount = qty * unit_price*
e. *defining default values for null fields on keys or values, e.g. if currency is null then "Euro"*
f. *standardising date formats, e.g. one source system contains dates the format yyyy-mm-dd and another one in dd-mm-yyyy*
g. *use logical operators for expressions and transformation formulas*

3) *Independence from underlying data source*
Board ETL rules are defined in a standard environment which is independent from the SQL dialect of the relational database of the different data sources, e.g. transform a date format from MSSQL and from an Oracle source using a single formula that doesn't depend on the T-SQL or PL-SQL syntax.

4.4   BOARD connector for SAP

SAP® Certified
Integration with SAP NetWeaver®

Harnessing data from SAP ERP is a notoriously difficult and sometimes a very expensive activity. To streamline this cumbersome process, BOARD offers a dedicated connector for SAP environments. The BOARD SAP connector was conceived to directly interoperate with all varieties of SAP interface technologies by means of 8 specialized components.

Easily installable and configurable, the BOARD SAP connector allows mapping SAP data into BOARD databases without writing any data extraction ABAP program.
A user-friendly graphical interface makes it easy to implement data flows in a secure and type-safe way, and to extract data to BOARD in few clicks with high performance and stability.

Tables

This method allows you to extract mass data directly from SAP tables. In order to optimize the network load, only the columns that have been explicitly selected are extracted

BAPIs and function modules

Using the Xtract BAPI component you can call BAPIs or RFC function modules directly from BOARD DataReaders

OHS (Open Hub Services)

Thanks to the capability to connect with SAP OHS, this component allows to be created, automatic extraction processes including different SAP objects (not only cubes, but also InfoObjects, ODS objects, texts, etc.)

Reports

Allows the use of existing ABAP reports and programs as data sources within the BOARD DataReader – making it possible to reuse existing logic and thus avoiding any reengineering. This is why this technique is often used in the area of FI/CO reporting

Queries

This component makes SAP queries directly available in BOARD. Thanks to the graphical user interface and the full metadata support, the result is available for further processing within seconds

BW-Cubes & BW-Queries

*This component extracts data sets from SAP BW InfoCubes.
Dimensions, key figures and attributes can be transferred by a simple drag and drop into the query output*

Hierarchies

This method allows users to extract any hierarchy type from SAP BW. If required, hierarchy texts can be extracted as well.

DeltaQ

DeltaQ component allows using the SAP BW Extractor API (the functionality that SAP BW relies on to get its data from the live system), making any business information immediately available – without the need to understand and work on the underlying table structure

4.4.1   How it works

The BOARD Connector for SAP is a middleware component to extract data from SAP systems and load it into BOARD. It has two components: Designer and Server.
The Designer allows the System Administrator to design and configure the extractions of SAP objects such as Tables, Views, Queries, BW-Cubes, BAPI function modules etc.
The Server SAP APIs (SAP Certified) text extracts data from SAP and passes it to BOARD

More specifically, the process can be split in these steps:

*1) **Extractions creation***
The System Administrator uses the BOARD Connector Designer to define a catalog of SAP objects (tables, queries, ABAP programs...) that can be exposed to BOARD: the Xtracts library.
The BOARD Connector connects to SAP through the RFCs, the standard SAP APIs.

*2) **Data Reader configuration***
The System Administrator maps the available extractors into BOARD databases
BOARD Data Reader connects to BOARD Connector through Web Services.

*3) **Extractor Execution***
BOARD SAP Connecter sends a request to SAP and passes the result set to BOARD.

5. PERFORMANCE AND SCALABILITY

5.1   The HBMP technology: the new in-memory frontier

At its core, the BOARD platform is based on a revolutionary proprietary in-memory data management technology named HBMP.

HBMP technology has been conceived with a precise and extremely ambitious goal: to exploit the full potential of in-memory computing (IMC) in terms of better performance, while improving the flexibility, scalability and ability to effectively support decision-making processes that have fueled international success for BOARD.

The result is a revolutionary solution that overcomes the limitations of standard BI in-memory technology, combining pure performance with the unique capability to support and digitalize simulation and planning processes across the enterprise.

The technology's name, Hybrid Bitwise Memory Pattern, reflects its two main differentiating characteristics: the Bitwise Memory Pattern and its unique Hybrid approach.

Bitwise Memory Pattern

The term "Bitwise Memory Pattern" describes BOARD's ability to store and map physical data through unique bits, or rather to "bitmap" them, by means of an innovative proprietary algorithm. In other words, BOARD maps the compressed multidimensional data structure (Pattern) at bit level (Bitwise) into the RAM (Memory). The new algorithm, designed to exploit all the potential of in-memory computing, brings enormous benefits, not only in terms of performance, but also in terms of data compression, parallelism and high-user concurrency. Furthermore, unlike what happens with the majority of in-memory solutions, the algorithm was conceived to ensure full bi-directionality (read-write) in the use of information. This allows full support of write-back and "on the fly" changes of the data structure, further improving these key BOARD points of strength.

Hybrid approach

Another fundamental characteristic of HBMP technology is its Hybrid nature, or rather the ability to manage data in three different ways:

1) **Full in memory:** *To do all processing against data held in-memory*

2) **On-demand in-memory:***To decide which subset of data to use in-memory as more appropriate to the workload (i.e. in-memory metadata for read-only operations; dynamic upload of physical data for write-back)*

3) **Hybrid in-memory:**
To store in-memory data indices and mapping and to leave data on the disk.

These three different ways of configuring the BOARD in-memory environment allow great flexibility in choosing the model that best suits the user's needs, both in terms of scalability and price-performance.

Any "out of memory" problem and related turnaround is eliminated, and at the same time BOARD offers several options to optimize the trade-off between performance and technological cost of ownership of BI and CPM implementations.

5.1.1   Speed, Concurrency, Data Volumes: the HBMP benefits

HBMP technology is not limited to the capability of storing data in-memory reducing the need for disk I/O.

It encompasses huge improvements in parallelism, data compression, and high user concurrency further enhancing BOARD's capability to handle large volumes of data and users at a high processing speed.

Parallel multicore processing

HBMP technology fully parallelizes BOARD processing to effectively leverage multi-core architecture. The effective spreading of the calculation load across CPU cores enables more effective workload management and a better handling of the concurrency.

Data Compression

Thanks to the adoption of a new mathematical model, the HBMP technology makes it possible to map in memory the multidimensional structure of the data at bit level. This new compression scheme considerably reduces the data space required to create cubes and makes the weight of BOARD databases very similar to the most advanced associated "read only" products, with the remarkable difference that BOARD databases are read and write and have the capability to distribute data changes along the hierarchies.

Concurrency

HBMP technology drastically improves high user concurrency in BOARD. Unlike traditional in-memory techniques that struggle or require massive amounts of RAM to support large numbers of concurrent users, HBMP provides the capability to support concurrent users sharing a unique server instance, thus minimizing the impact of concurrent users on memory usage.

5.2   HBMP: the architectural advantage

Behind the important benefits offered by HBMP technology lie three fundamental architectural factors that make it superior to any other in-memory solution available today.

Bi-directional algorithm

The mathematical model according to which the data is indexed and mapped in-memory has been designed and built to support not only the output process (read), but also the logically inverse input processes (write).

In-memory metadata

HBMP technology does not just allow for in-memory data management, but also ensures the ability to run the meta-data processing in RAM. From a practical point of view this ability makes it possible to change the database structure and mapping and to see the changes reflected in applications and analyses "on the fly" without having to reload the physical data. Moreover, it is from the combination of the in-memory metadata processing and the bi-directional algorithm that the HBMP technology derives its Hybrid nature, i.e. the ability to administer the mapping and indexing of data in-memory, deciding whether to physically store it in RAM or on disk.

Single server instance

HBMP technology is characterized by the ability to support concurrent users with a single server instance. This capability makes the impact of high concurrent users on the RAM resources almost insignificant, totally differentiating HBMP from all the main in-memory technologies, which are heavily impacted.

5.3   In-memory cluster: Horizontal Scalability

BOARD offers an advanced In-Memory Server Cluster architecture making it possible to share BOARD server workload on a virtually unlimited number of nodes.

This functionality ensures horizontal scalability with read and write support, and provides cutting-edge technological resolutions to four main architectural issues usually experienced in large and geographically distributed implementations, namely:

User Scalability

When the number of users becomes so large that it affects the performance of a single server, BOARD Cluster offers the capability to automatically allocate users to the available nodes through a load balancing mechanism.

Data Partitioning

The management of Big Data projects often requires handling and analysing very large data volumes that affects performance of a single server. In this case, BOARD Cluster can be utilized to partition data on different nodes while maintaining a consolidated "Single Point of Data". Through an automatic sharing process, the BOARD database is transparently partitioned across nodes, allowing it to scale-out read-write Business Intelligence and Performance Management applications without requiring any changes to them.

Process scalability and 24/7 availability

The deployment of large worldwide solutions with users working around the clock and in different time-zones requires the solution to be available at all times leaving no idle time window for performing typical overnight operations such as data feeds and bulk loads from source systems or bulk loads.

In this case, the BOARD Cluster allows the system architect to distribute the execution of bulk loads or other mass processes across different nodes of the cluster. This not only optimizes the overall system performance by scaling out processes on different servers but it also isolates users from the impact of running mass data processes.

Geographic Latency

Performance of worldwide or long-distance projects are often affected by a problem known as "Geographic Latency" or, more simply, infrastructural deficiencies in the network connection. BOARD Cluster helps companies minimize these kinds of problems, making it possible to handle local data, users and operations on local nodes while having all the outcomes consolidated in a "Single Point of Data" through a low-impact and super-fast in-memory update process.

This way, BOARD can ensure fast read and write performance to decentralized users, without impacting applications design and, more generally, the overall functioning of the system.

BOARD In-memory cluster: architectural schema

CLIENTS   MASTER
                  NODE   SLAVE
                                NODES
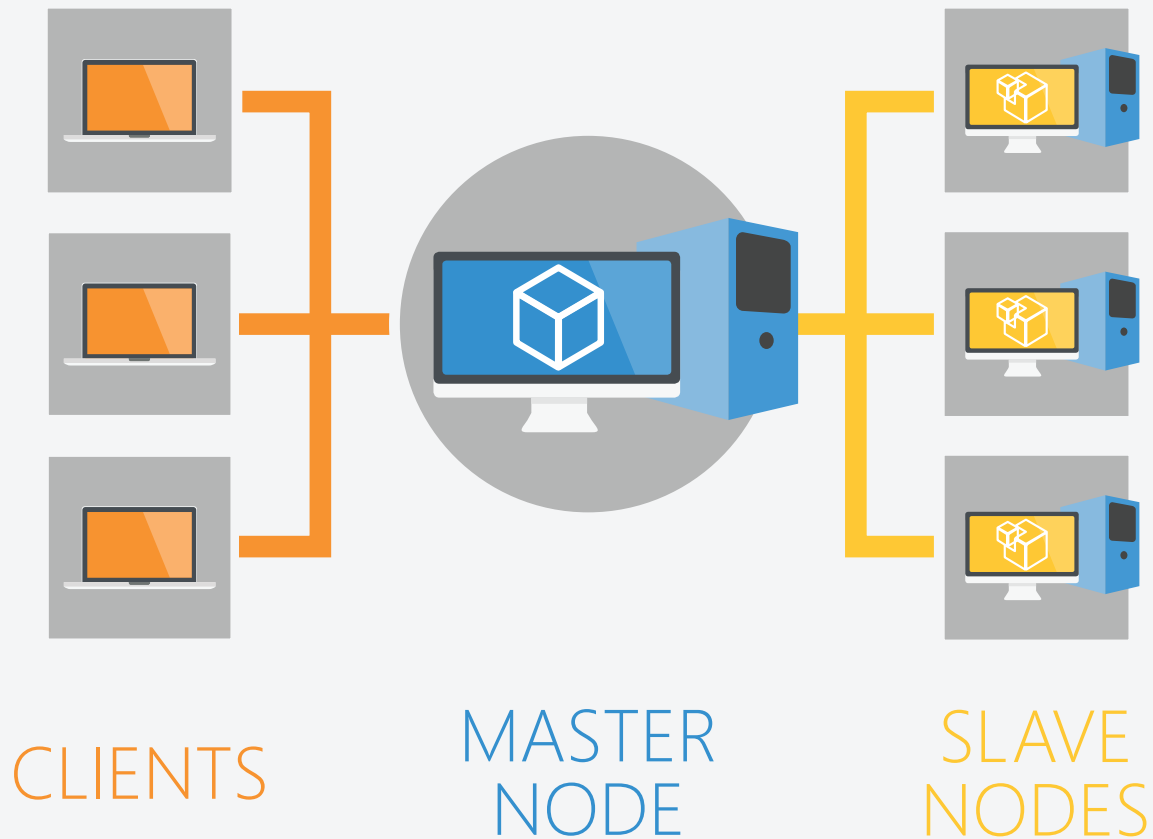
5.3.1   How the BOARD Server Cluster works:

The BOARD Cluster is based on a hub-spoke model where a Master Node serves as a single point of access for a dynamic number of Slave Nodes.

The Master Node addresses users, processes and data to slaves based on the policy implemented by the system administrator (see the paragraph "Flexible Load Balancing Logics"). The Slave nodes work locally on their In-Ram databases and send delta updates of data to the Master node. The Master node collects and merges data received sequentially (last one wins) and, after merging, pushes updated data to all slaves.

At a given point in time, the data on a Slave node can be different (more recent) to the data in another Slave node, but eventually all the Slave nodes become consistent (Eventually Consistent Model). The whole process is performed in RAM using the BOARD HBMP In-Memory technology, thus making it extremely efficient: each single node can take charge of its own data fetch and calculation processes (the most time-consuming activities) while only the delta data outcomes are synchronised between servers through a super-fast in memory update.

5.3.2   Adaptive Load Balancing

BOARD In-memory Cluster can be configured in several ways, to perfectly adapt to different use cases:

1) Users: *BOARD Master Server automatically redirects users to the various slave nodes. This approach is suited to handle a large amount of users when there are no other specific architectural requirements.*

2) Security Profiles: *BOARD Master Server redirects users to a particular slave node based on their security profiles. This approach is utilized to manage different kinds of operations on different nodes (i.e. delegate to a single server all the data-reading processes) or to handle geo-clustering (i.e. addressing all the users connecting from a specific zone to a local server).*

3) Performance Index: *BOARD Master Server redirects users on a particular slave node based on a Performance Index, defined by the system administrator. This approach is utilized to allow the system administrator to fine-tune users' workload based on the performance of the underlying HW architecture.*

4) Ad-hoc configuration (Specialized Nodes): *When managing geographical distributed, complex planning and simulation processes, an automatic load-balancing configuration is often an ineffective architectural choice. The system administrator needs to have the ability to define which data, users and processes to handle on any single node in order to align the system architecture to the specific functional requirements, resolve potential workload bottlenecks and optimize the performance of the whole process.*
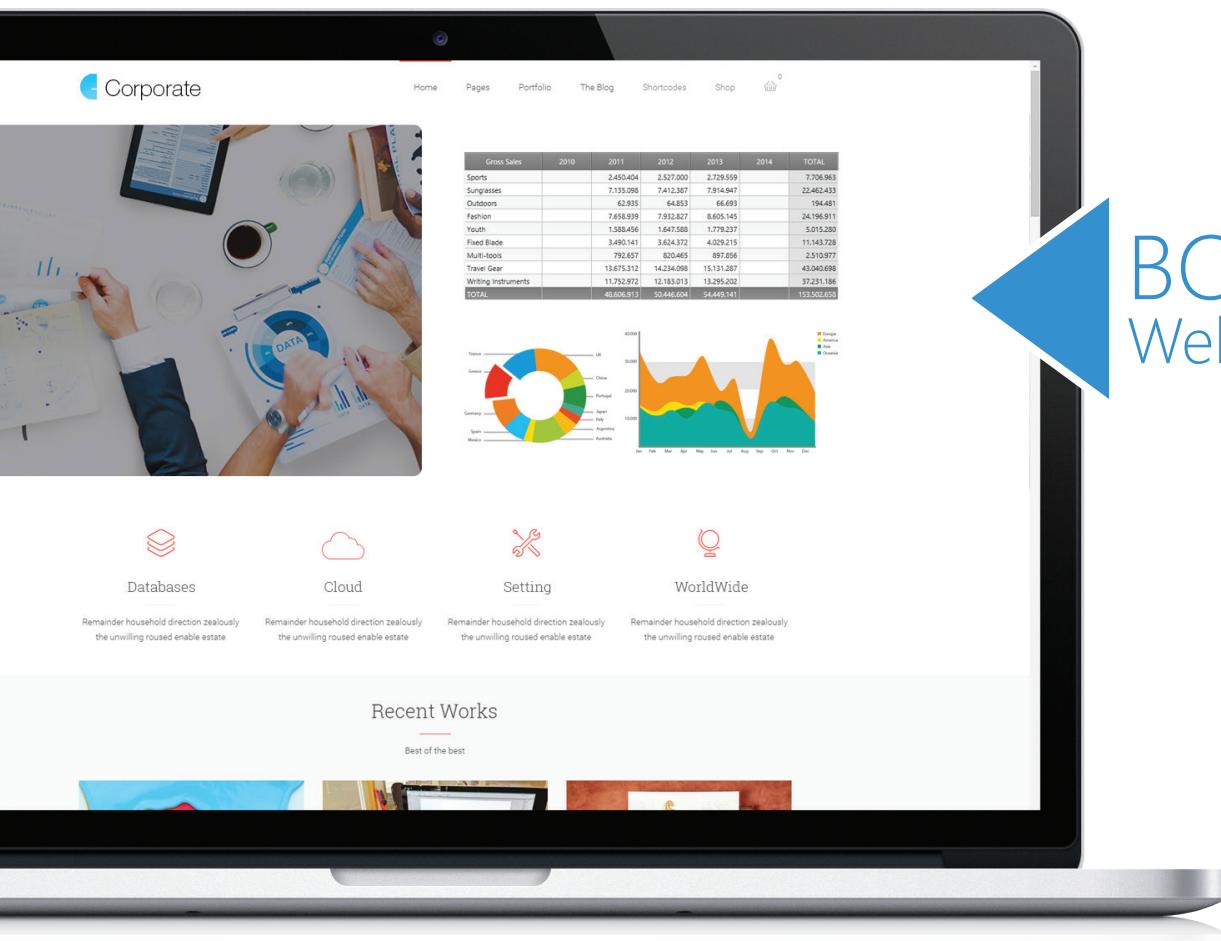
6. INTEGRATION AND PORTAL INTEGRATION

6.1 Web Mash-up and portal integration

The BOARD Web Client is based on Microsoft Silverlight.

The BOARD Web Client is a browser-based runtime environment which can run BOARD applications directly on the web.

Through Java Script it is possible to call the methods exposed by the Silverlight client and interact with BOARD applications.

These capabilities combined with the Single Sign On (see chapter 3), enables BOARD to be easily integrated with any company portal.

6.2   Sharepoint integration

BOARD Web and SharePoint portal are both based on SOA architecture, Microsoft .Net Framework and Silverlight technology.
This full technologic alignment makes the integration between the two systems quick and easy.
In particular, the integration can be handled at three different levels:

Web Page Viewer

The simpler option to visualize a BOARD web screen on a Web Part Page is to use the Page Viewer. The Page Viewer is a default Web Part in SharePoint server and SharePoint foundation. The SharePoint admin simply enters a hyperlink, file path or folder name to link to the content. The linked content of the Page Viewer Web Part is isolated from other contents on the Web Part Page by using an HTML IFRAME element.
This element ensures that any HTML elements that are displayed as content in the Page Viewer Web Part do not conflict with other HTML elements on the Web Part Page.
The Page Viewer Web Part displays content asynchronously from the rest of the page, disconnecting the page performance to the linked object performance. This means that users can view and utilize other Web Parts on the page, even if the link happens to take a long time to return the content.
This solution allows users to interact directly with the BOARD screen, but it is not possible to link the BOARD environment to any SharePoint object, nor to define any access parameter (BOARD application is consumed "as is").

BOARD Silverlight Client Web Parts

Separately Silverlight and SharePoint perform their tasks well, but in today's web integration environment each can benefit from integration with the other application.
SharePoint 2010 has built-in support for Silverlight Web Parts making it easy to get Silverlight applications up and running.
Web parts in SharePoint allow developers to create "composite-able" units of functionality that power users can then put together to build SharePoint sites and applications.
The BOARD Silverlight Client Web part is extremely easy to use. A Silverlight Web Part makes it possible to embed a BOARD application into a SharePoint site simply by adding the Web Part itself to the page and setting the path to the BOARD Web server.
This solution is based on the configuration of a turn-key component so it is extremely simple to implement and makes it possible to define access parameters to the BOARD application (Capsule, Screen, Menu visualization, etc.)

Html Page Integration

SharePoint allows users to visualize and directly access html pages.
Through a traditional HTML Editor or Microsoft SharePoint Designer it is possible to develop a web page embedding the html code needed to call the Silverlight client where BOARD runs. This mashup solution allows the dynamic interaction between the BOARD Web Client and the web page. The html page can be a created in a SharePoint library or in an external website and then retrieved from there.
This solution permits a comprehensive integration between BOARD and the SharePoint environment but will require the use of HTML and Java Script programming.

6.3   BOARD Web Services

BOARD Web Services allow flexible integration of BOARD with external applications.
The BOARD server, conceived to be used as a Web Service, manages SOAP messages (usually formatted in XML) and returns XML data structures.
The services supplied by the BOARD server can be used by other Web Services, by legacy systems, by Java, .NET and web applications, by company portals and, more generally, by any application, in a simple and standardized way (W3C Standard).
By providing these services, BOARD makes it possible to extend advanced Business Intelligence and Performance Management functions to the existing company application environment.
BOARD WebServices are essentially based on two groups of methods:

> **1) Data dictionary functions:** *(database metadata browser) that provide the list of Cubes (Measures), Dimensions, Hierarchies and Members*
>
> **2) Query execution functions:** *that allow access to data*