

BOARD Technology Overview



Better decisions. Better business.

CONTENTS

- 1 PLATFORM OVERVIEW.....6
 - 1.1 Analyze, predict, plan and simulate in a single platform.....6
 - 1.2 Self-service analysis meets data governance: a two-way process.....8
 - 1.3 Product architecture.....10
- 2 DATA INTEGRATION.....14
 - 2.1 Data federation.....16
 - 2.2 BOARD Connectors.....17
 - 2.3 Multidimensional data sources.....17
 - 2.4 BOARD Connector for SAP.....18
 - 2.5 Essential ETL functionalities.....20
 - 2.6 Data Fast Track.....21
- 3 BOARD ENGINE.....22
 - 3.1 Databases.....22
 - 3.1.1 The HBMP technology: the new in-memory frontier.....23
 - 3.1.1.1 HBMP benefits: speed, concurrency and data volumes...24
 - 3.1.1.2 HBMP: the architectural advantage.....24
 - 3.2 Business logic.....25
 - 3.2.1 Defining the business logic.....25
 - 3.2.2 Predictive and advanced analytics (BEAM).....26
 - 3.3 Application.....28
 - 3.3.1 The toolkit: the fastest way to build analytic applications.....28
 - 3.3.2 Storytelling.....29
 - 3.3.3 In-context Collaboration.....30
 - 3.3.4 Cognitive technologies.....25

- 4 CLIENTS.....33
 - 4.1 BOARD Web.....33
 - 4.2 BOARD Mobile.....34
 - 4.3 BOARD Office Add-Ins.....35
 - 4.4 BOARD Windows Client.....36
- 5 SECURITY.....37
 - 5.1 Multi-tier security model.....37
 - 5.2 Authentication methods.....38
 - 5.3 Role based authorization.....39
 - 5.4 Application authorization.....39
 - 5.5 Data access.....39
 - 5.6 Advanced security management.....40
- 6 INFRASTRUCTURE.....41
 - 6.1 Cloud or on-premises.....41
 - 6.2 In-memory cluster: horizontal scalability.....42
 - 6.2.1 How the BOARD Cluster works.....44
 - 6.2.2 Adaptive load balancing.....44
 - 6.3 Interoperability and portal integration.....45
 - 6.3.1 Web mash-up and portal integration.....45
 - 6.3.2 SharePoint integration.....45
 - 6.3.3 BOARD web services.....46

1. OVERVIEW

1.1 Analyze, predict, plan and simulate in a single platform

BOARD is a unified platform for decision-making that combines business intelligence, advanced analytics, planning and simulation capabilities into a single environment. These capabilities work in concert to provide the fastest and easiest way to build highly customized analytical and planning applications.

The all-in-one approach means that BOARD provides all the capabilities needed to build, manage and maintain analytic solutions for:

- *Reporting and interactive dashboarding*
- *Ad-hoc analysis and data discovery*
- *Predictive analytics and simulation*
- *Budgeting, planning and forecasting*
- *Profitability modeling and financial consolidation*
- *Scorecarding and strategy management*

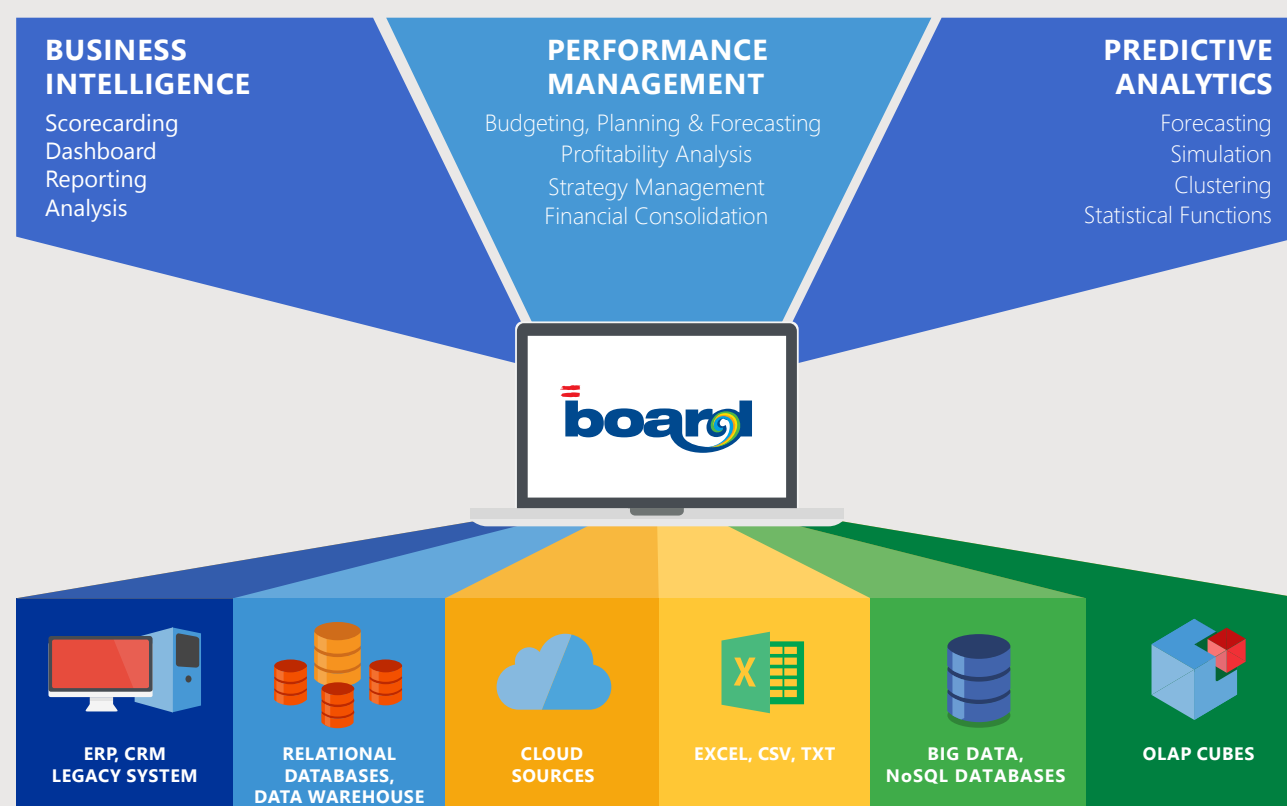
BOARD merges data analysis and simulation into the planning, workflow and controlling cycles of organizations; transforming data into insights, insights into predictions, predictions into plans and plans into traceable and monitored actions at a strategic, financial and operational level.

All solutions created in BOARD share the following elements:

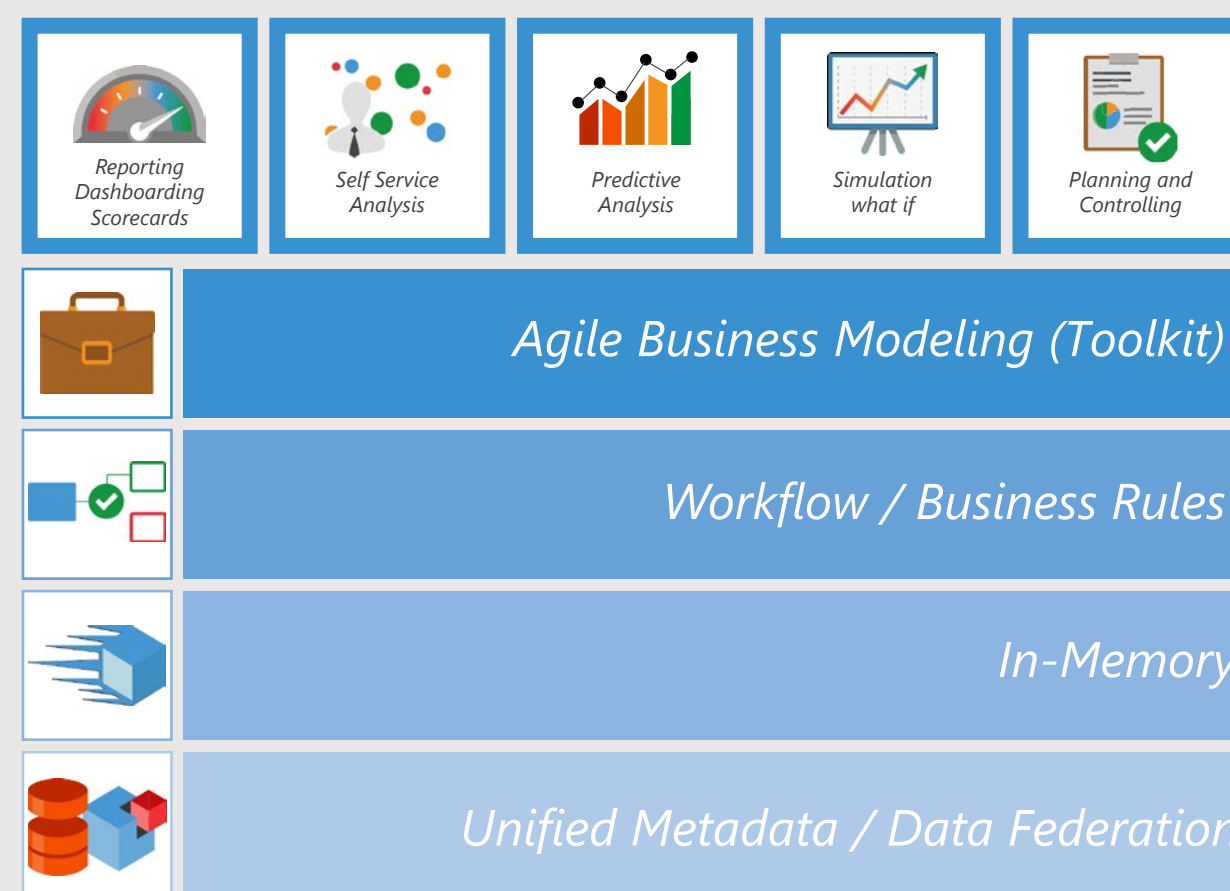
- *Single logical view of data*
- *Single and shared metadata across all types of solutions (BI, Planning, Simulation, Analytics)*
- *Single user interface*
- *Single security*
- *Single administrative environment*
- *Single visual modeling environment*

Integrating these elements into a unified platform creates tremendous technical efficiencies, which BOARD leverages to provide an unrivaled combination of speed, ease of use, flexibility and minimal cost of ownership. This allows BOARD customers to build and maintain their own bespoke solutions at a fraction of the time and cost required for traditional purpose-built applications.

BOARD All-in-One approach



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



1.2 Self-service analysis meets data governance: a two-way process

Self-service business intelligence provides valuable benefits in terms of agility, velocity and user independence. However, if not combined with a proper data governance process, it introduces many long-lasting problems, including: an environment of proliferating BI silos, inconsistent business definitions, no data lineage and no single version of the truth.

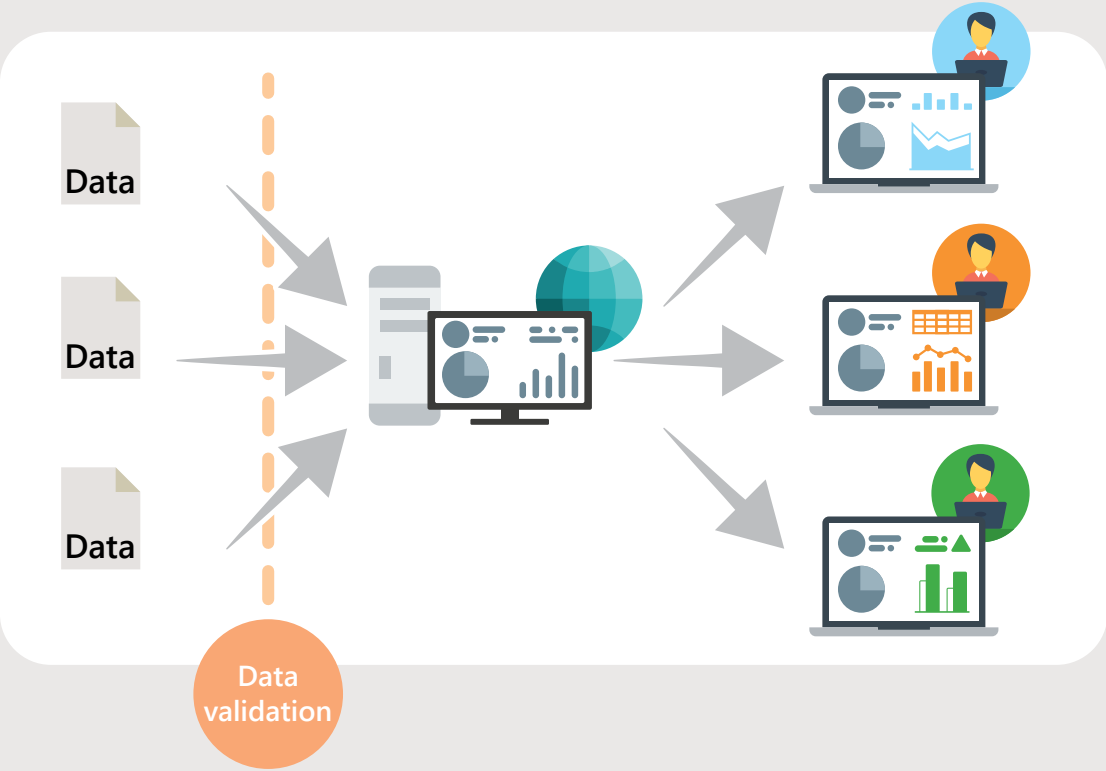
BOARD offers an innovative approach to governed data discovery, empowering business users with full self-service capabilities while allowing the organization to maintain full control of data security, integrity and the degree of self-reliance.

From one side, business users can customize any number of centrally-authored reports and visualizations and then save them to a personal workspace. This way, users can create their own personal perspective on enterprise reports, underpinned by a validated and centralized system of record.


From the other side, users can import their own data into a sandbox environment and use it for independent analysis. If the insights provided by the new data are relevant to others, the analysis and the underlying data can be easily posted to a common repository and, once validated by an administrator, promoted and shared across the organization.

This two way-process fully supports the self-service needs of the organization, in a trusted and validated environment, backed by the capability to trace the data lineage back to its source and to control how the data was manipulated before being analyzed.

*From the center to the Edges:
Users create their own customized version of corporate reporting*



*From the Edges to the center:
Users created their own reports that can be validate and included into corporate reporting/environment*



1.3 Product architecture

The BOARD platform ensures all the scalability, security and data-governance capabilities that IT needs to deploy and support enterprise-class business intelligence and planning projects. BOARD's multi-tier architecture offers:

- **Scalability**
BOARD been designed to ensure high levels of scalability in terms of concurrent user requests, data volumes and the scope/complexity of applications. Its integrated in-memory clustering technology supports large (thousands of users) global deployments and avoids problems with geographic latency.
- **Performance**
BOARD offers outstanding performance, not only in terms of data aggregation and presentation, but also in handling large and complex planning processes with thousands of users performing concurrent data-entry. This is made possible by the combination of BOARD's hybrid in-memory technology (HBMP) with a powerful and efficient multidimensional database that includes sophisticated sparse data management technology.
- **Consistent experience**
Any application built in BOARD can be accessed by multiple clients (web browser, BOARD Windows Client, mobile apps, Office add-in) providing a consistent user experience across all devices.
- **Robust security**
Security is pervasive across the whole platform, from authentication to cell level authorization. BOARD offers a robust, comprehensive and easy to manage enterprise-class security system.
- **24 x 7 availability**
The flexibility of BOARD's data model allows users to make changes in real time and to perform online data updates, which makes it possible to provide organizations with 24 x 7 availability.

BOARD's functional architecture includes the following elements:

Data integration

From relational databases to cloud applications, from big data to cloud data stores, it does not matter which data you access, nor where it is stored. BOARD offers a suite of connectivity tools to streamline access to virtually any data source.

BOARD Engine

The BOARD Engine encompasses all the functionality required to model, build and manage solutions on the platform and is based on three building blocks: databases, business logic and applications.

- **Databases:**
Databases provide the capability to access, normalize, federate and manage data.
- **Business logic:**
A comprehensive set of functionalities and business rules to implement, run and manage analysis, planning, budgeting, forecasting, consolidation, allocation and simulation logic – seamlessly linking data with business processes.
- **Applications:**
The end-user environment that allows users to view, interact with and analyze data. Non-technical users can become developers, quickly building applications without coding, leveraging all the capabilities offered by the databases and the business logic.

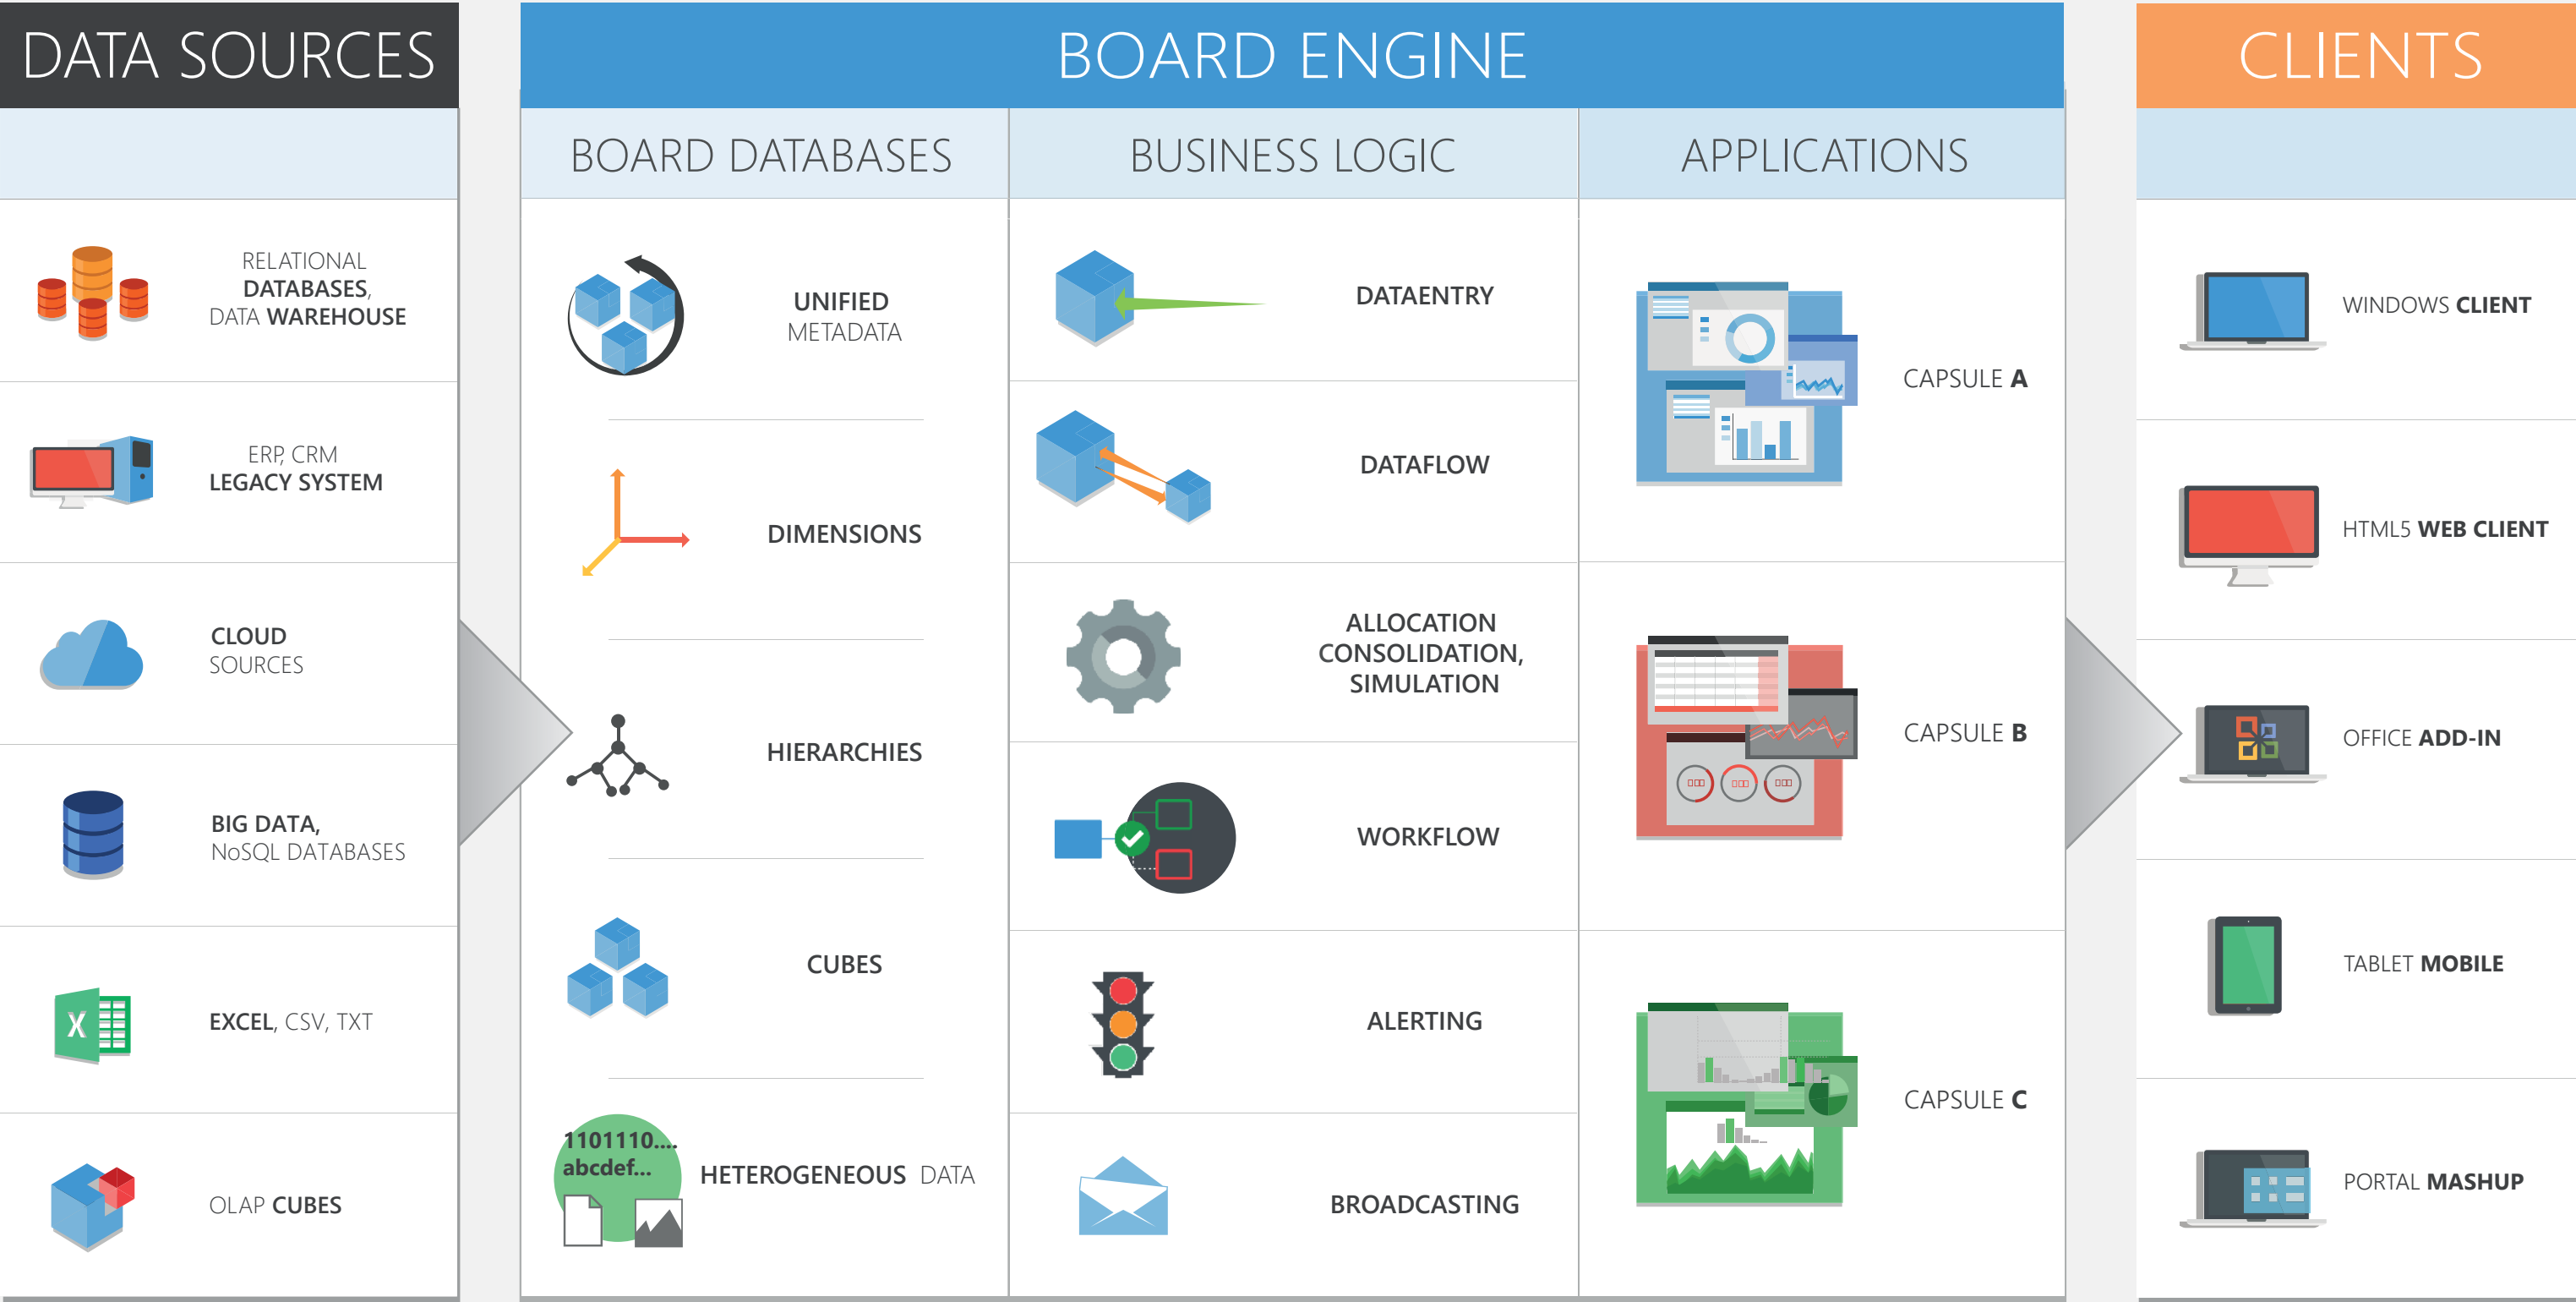
Clients

BOARD ensures a consistent experience across web, mobile and desktop environments so developers can build a single application and deploy it to users on variety of devices. BOARD applications can be accessed through the following clients:

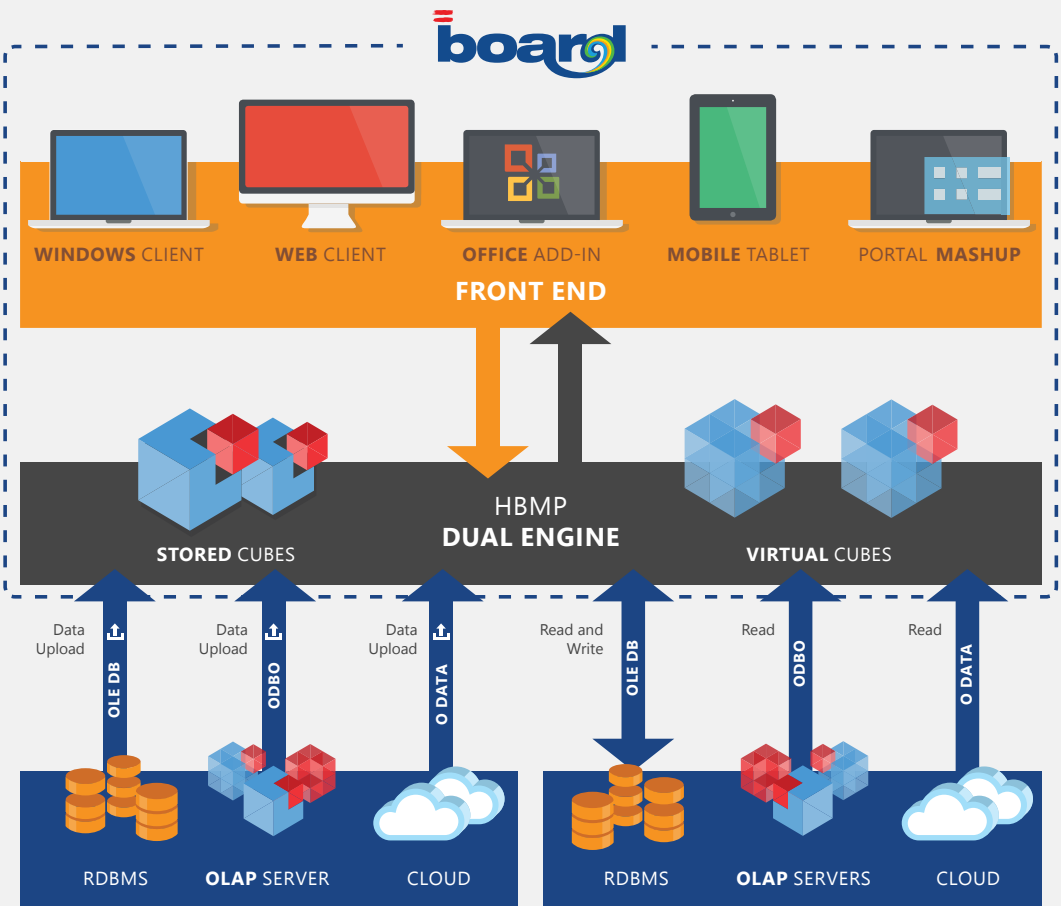
- **BOARD Web (HTML5)**
- **BOARD Mobile**
- **BOARD Office Add-Ins**
- **BOARD Windows Client**

Infrastructure & Security

BOARD offers an enterprise-grade infrastructure to support the largest worldwide deployments. This includes strong and flexible multi-tier security, custom-tailored cloud and on-premises networking, as well as advanced in-memory clustering capabilities.



2. DATA INTEGRATION



While maintaining data within a single enterprise data warehouse can provide substantial benefits to an organization, business complexity often makes this too demanding, difficult or even impossible. In practice, large volumes of data are frequently stored in a wide variety of sources and platforms, both internal and external to the enterprise.

The BOARD platform provides the capability to connect and integrate data across virtually any type of data source, including:

- *Relational databases and data warehouses (e.g. Oracle, SQL Server)*
- *Enterprise applications (e.g. SAP ERP)*
- *Cloud based sources (e.g. Salesforce, Google Analytics)*
- *Column oriented databases (e.g. Teradata, SAP HANA)*
- *Big Data/NoSQL databases (e.g. Hadoop, Impala)*
- *Multidimensional sources (e.g. SAP BW)*
- *Web services*
- *Excel, CSV and TXT files*

Through its physical and logical multidimensional database, BOARD offers a single metadata gateway that makes it possible to use data as if it were native to BOARD. Data is normalized and organized within BOARD, allowing end-users to read, write and update the data, regardless of the source.

Behind the scenes, the administrators have the freedom to decide which data to store within BOARD's multidimensional database (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 users with a single and seamless vision of the data.

BOARD also supports write back to both its multidimensional database as well as to relational data sources, making it possible to integrate performance management processes with enterprise applications (e.g. storing forecasting/planning data directly within the ERP).

Access to these data sources is provided through the following technologies:

- *Native pre-built data connectors*
- *Open Database Connectivity (ODBC) and OLE DB to connect with relational databases*
- *ODATA, JSON and REST to connect with cloud based sources*
- *OLE DB for OLAP (ODBO) to connect with multidimensional databases*
- *BOARD connector for SAP*
- *Native text file reader*

2.1 Data federation

Data federation is the process of gathering data from disparate data sources without transferring the original data. Rather than combine all of the source data into a new database, physical data is left in place and metadata - data that describes the data and its location - is utilized to create a virtual database.

BOARD's data federation capabilities (ROLAP engine) provide the ability to read and write data stored in a RDBMS as if it were stored in the BOARD multidimensional database. More precisely, metadata is stored in BOARD's virtual cubes; whenever a user views a report based on these cubes, a connection to the external relational database is established and data is retrieved from the relational table through a SQL query.

Anytime a user runs a request that accesses multiple relational sources, BOARD automatically: Generates a series of queries for each database needed by the report/analysis, connects to the corresponding data sources for execution, feeds the BOARD virtual cubes with the final results from each of the queries and transparently makes the final merged results available in the report/analysis

From an architectural standpoint, this capability allows developers:
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 ability to access this data. 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 ability to access the lowest level of detail through virtual cubes whenever needed.
To create near real-time analytic models by making RDBMS data directly available in the BOARD MOLAP environment.

2.2 BOARD Connectors

BOARD offers a comprehensive set of pre-built dedicated data connectors that minimize the time and effort required to access and use your data. From relational databases to cloud applications, from big data to cloud data stores, it does not matter which data you tap, nor where it is stored.

Big Data	Relational and Analytics	SAAS/Cloud
Amazon EMR Hive Apache Cassandra Apache Hadoop Hive Cloudera CDH Hive Cloudera Impala DataStax Enterprise Hortonworks Hive IBM BigInsights Hive MapR Hive MongoDB Pivotal HAWQ Pivotal HD Hive Spark SQL	Amazon Redshift IBM DB2 IBM Informix Microsoft SQL Server MySQL Community MySQL Enterprise Oracle Database Pivotal Greenplum PostgreSQL Progress® OpenEdge® SAP Sybase ASE SAP Sybase IQ Teradata	Financial Force Force.com Applications Google Analytics Hubspot Marketo Microsoft Dynamics CRM Microsoft SQL Azure Oracle Eloqua Oracle Sales Cloud Oracle Service Cloud Progress® Rollbase® Salesforce.com ServiceMAX SugarCRM Veeva CRM

2.3 Multidimensional data sources

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

- **Microsoft Analysis Service**
- **Oracle/ Hyperion Essbase**
- **SAP BW**
- **IBM Cognos TM1**

ODBO is a Microsoft specification and an industry standard for multidimensional data processing specifically designed to access multidimensional data sources. XMLA is essentially a newer standard interface. Both of these technologies operate using the Multidimensional Expressions (MDX) language to query cubes.

2.4 BOARD Connector for SAP



Harnessing data from SAP ERP is a notoriously difficult and sometimes a very expensive activity. BOARD offers a dedicated connector for SAP environments to streamline this process. Easily installable and configurable, the BOARD SAP Connector allows mapping SAP data into BOARD databases without writing an ABAP program for data extraction.

A user-friendly graphical interface makes it easy to implement data flows to extract and load data into BOARD in a safe and secure manner, all with high performance and stability.

The BOARD SAP connector interoperates with all varieties of SAP interfaces by means of 8 specialized components.



Tables

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



Queries

This component makes SAP queries directly available in BOARD



BAPIs and function modules

Using the Xtract BAPI component, you can call BAPIs or RFC Function Modules directly from BOARD Data Readers



BW-Cubes & BW-Queries

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



OHS (Open Hub Services)

Connect with SAP OHS. This component allows automatic extraction processes to be created, including to different SAP objects (not only cubes, but InfoObjects, ODS objects, texts, etc.).



Hierarchies

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



Reports

Use existing ABAP reports and programs as data sources within the BOARD Data Reader, making it possible to reuse existing logic. This technique is often used in the area of FI/CO reporting



DeltaQ

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

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 a system administrator to design and configure the extractions of SAP objects such as tables, views, queries, BW Info 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 into these steps:

- **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 in the Xtracts library. The BOARD Connector connects to SAP through the RFCs, the standard SAP APIs.
- **Data Reader configuration**
The System administrator maps the available extractors to BOARD databases. BOARD Data Reader connects to BOARD Connector through Web Services.
- **Extractor Execution**
BOARD SAP Connector sends a request to SAP and passes the result set to BOARD.

2.5 Essential ETL functionalities

While BOARD is not a dedicated ETL product, it does offer capabilities for diverse source systems to directly feed BOARD databases without the need for intermediate data staging layers. This is unique when compared to other BI products, which typically require the source data to be cleansed and organized into either a star or snowflake schema, which is costly, rigid and time consuming.

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

- **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)
- **Normalization and conversion**
 - a. Translating coded values (e.g. source system contains 1 for male and 2 for female but we want to use M for male and F for female)
 - b. Encoding freeform values (e.g. mapping DE to Germany, FR to France)
 - c. Standardizing codes from multiple sources (e.g. one source system uses 01 for a given product and another system uses ABC)
 - d. Deriving a new calculated value (e.g. $sales_amount = qty * unit_price$)
 - e. Defining default values for null fields (e.g. if currency is null then "Euro")
 - f. Standardizing date formats (e.g. one source system contains dates in a format of yyyy-mm-dd and another one in dd-mm-yyyy)
 - g. Use logical operators for expressions and transformation formulas
- **Independence from the underlying datasource**
Board ETL rules are defined in an independent environment from the SQL used to query the relational databases of the different data sources, e.g. transform a date format from a MSSQL source and from an Oracle source using a single formula that doesn't depend on the T-SQL or PL-SQL syntax.

2.6 Data Fast Track

The Data Fast Track feature is for self-service data modelling, mash-ups and uploading. With Data Fast Track, business users can quickly create compelling visual analyses and analytical applications from raw data in just a few clicks, without any assistance from IT.

Users connect to a dataset and BOARD will automatically build a data model that reflects the relationships inherent in the source data. Users can then enrich the data model by blending additional datasets or by creating derived fields in an Excel-like environment.

Once the data model is finalized, users can immediately create their analyses using BOARD's drag-and-drop environment.

- **Allows business users to create analysis from raw data**, in just a few clicks.
- **Smart data blending** from different sources. Automatically merge diverse data sources into a single data model.
- **Automatic discovery and classification of data types** (measures, dimensions, hierarchies and entities).
- **Self-service data manipulation.** A spreadsheet-like environment allows user to create new measures with formulas, to cut and paste data and to edit meta-data mappings.

3. BOARD ENGINE

The BOARD Engine encompasses all the functionality required to model, build and manage solutions on the platform and is based on three building blocks: databases, business logic and applications.

- Databases**
Databases provide the capability to access, normalize, federate and manage data
- Business logic**
A comprehensive set of functionalities and business rules to implement, run and manage analysis, planning, budgeting, forecasting, consolidation, allocation and simulation logic – seamlessly linking data with business processes.
- Applications**
The BOARD end-user environment that allows users to view, interact with and analyze data. Developers can quickly build and run applications without coding, leveraging all the capabilities offered by the databases and the business logic.

The BOARD Engine 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 Engine also performs user authentication, applies security restrictions or privileges and then dispatches user requests as separate execution threads to carry out the analytical processing.

Communication between the clients and the engine use a proprietary protocol named ROAR (Remote Object Access & Replication) which provides extremely efficient communications and offers high performance over low bandwidth connections.

3.1 Databases

BOARD provides a single view of organizational data through its combination of physical and logical multidimensional databases. BOARD acts as a metadata gateway that allows all source data to be used as if it was native to BOARD. Data is normalized and organized in BOARD data models, allowing users to perform reads, writes and updates regardless of data sources.

Behind the scenes, administrators have the freedom to decide which data to store into BOARD's multidimensional databases (MOLAP), and which to leave in place and manage through BOARD data federation capabilities (ROLAP). BOARD's DUAL Engine automatically manages the two different options, providing users with a single and seamless vision of the data.

This model also applies to data-entry processes; BOARD supports write-back not only on its data model, but also directly to relational data sources.

3.1.1 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 brings enormous benefits, not only in terms of performance but also in terms of data compression, parallelism and high-user concurrency. Unlike most in-memory solutions, the algorithm ensures full bi-directionality, which supports write-back and "on the fly" changes of the data structure.

Hybrid approach

The hybrid nature provides the ability to manage data in three different ways:

- Full in memory:**
All processing is performed against data held in-memory
- On-demand in-memory:**
Determine which subset of data to use in-memory and which to store on disk. E.g in-memory metadata for read-only operations, dynamic upload of physical data for write-back
- Hybrid in-memory:**
Use the data indices and mapping in-memory but store the rest of the data on disk

These three different ways of configuring the BOARD in-memory environment allow great flexibility in choosing the model that best suits the solution requirements and to make optimal tradeoffs between scalability, price and performance.

3.1.1.1 HBMP benefits: speed, concurrency and data volumes

The HBMP technology is not limited to just storing data in-memory to reduce disk I/O, it also provides huge improvements in parallelism, data compression and user concurrency. In functional terms, this allows users to perform analysis and build simulation models on large volumes of highly granular operational data.

Parallel Multicore Processing

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

Data Compression

The HBMP technology makes it possible to map the multidimensional structure of the data in memory at bit level. This compression scheme considerably reduces the space required to store and manage data, making the “weight” of BOARD databases similar to the most advanced read-only products.

Concurrency

HBMP technology supports high levels of user concurrency, unlike traditional in-memory techniques that struggle with or require massive amounts of RAM to support large numbers of concurrent users.

3.1.1.2 HBMP: the architectural advantage

Three fundamental architectural factors that make it superior to other in-memory solutions in the market.

Bi-directional algorithm

The mathematical model by which the data is indexed and mapped in-memory supports both read and write operations and can distribute data changes along hierarchies.

In-memory metadata

HBMP technology also performs the meta-data processing in-memory, not just data management like with other products. This allows the database structures and mappings to be changed “on the fly” and to immediately see the changes reflected in the applications without having to reload the physical data. The combination of the in-memory metadata processing and the bi-directional algorithm is where HBMP technology derives its hybrid nature, i.e. the ability to administer the mapping/indexing of data in-memory and deciding whether to physically store it in RAM or on disk.

Single server instance

HBMP technology supports concurrent users sharing a unique server instance. This minimizes the impact of high concurrency on memory resources and is dramatically different than other in-memory technologies, which are heavily impacted by concurrent user requests

3.2 Business logic

BOARD provides a comprehensive set of functionalities to implement the business logic that organizations need to support their analytic and planning processes.

Such processes include driver-based allocations, eliminations, conversions, consolidations, initializations, top down and bottom-up reconciliations and multidimensional calculations. More generally, this encompasses all the necessary operations to implement company-wide planning, forecasting and simulation applications and includes functionality like workflow, scenario management and versioning.

3.2.1 Defining the business logic

The business logic is captured within several different areas of the platform, depending on the desired functionality. The most common are:

- Procedures
- Report and visualization objects within application screens
- BOARD’s predictive engine (BEAM)

Regardless of where the logic is captured, it is always defined using functional business terms, not code, within an agile and interactive modeling environment (the “toolkit”). This allows the business users themselves to quickly build customized solutions that perfectly fit the organization’s requirements. The agile nature of the “toolkit” means that solutions can also be quickly adapted to ever changing business needs.

It’s tempting to think that since the logic is easy to define and doesn’t require coding, then it may lack the flexibility to support complex, real-world use cases. In fact, the opposite is true. BOARD users define the business logic using an intuitive, parameter-based approach. This is not only easier to learn but since the business logic is tightly coupled with the database in an integrated system (i.e. “data-aware”), users can make fundamental changes to the structure and contents of the underlying database and have these changes automatically propagated into the business logic.

Compare this with a traditional, code-based solution where the manual maintenance required to keep the hard-coded business logic in sync with the underlying data severely limits the scope and complexity of the solutions that can be provided.

While the business logic is defined using functional business terms, the logical operations that the underlying engine supports are highly advanced.

Some examples:

- *User-defined formulas, which can leverage over 325 purpose-built functions*
- *Native sequencing of dependent formulas*
- *Mixed-dimensional and mixed-grain operations*
- *Contextually-aware lookups to any level of detail, whether stored or calculated*
- *Aggregation rules (e.g. sum, avg, max, min, count)*
- *Layered filtering*
- *Procedure modeling*
- *Predictive analytics*
- *Clustering*

3.2.2 Predictive and advanced analytics (BEAM)

You can embed advanced analytical and predictive capabilities into your business processes with the agility necessary to operate at the ultra-rapid pace of today's business. And do it in a simple way, allowing users to take the lead, without having to depend on an army of experts like with the traditional data-mining and analytics products.

BOARD BEAM offers a groundbreaking solution to incorporate all the power of predictive analytics into daily business operations and the decision-making process, making it easier than ever for anyone to quickly gain powerful business insights and take action.

Turn your history into business prediction

Exploit all the power of BOARD BEAM's automated predictive modeling to execute extremely accurate forecasts in a fraction of the time and effort compared to traditional analytic solutions.

Thanks to its automated modeling approach, BOARD BEAM can automatically evaluate the characteristics of each time series, producing a suitable model, and can run forecasts thousands of times faster than competing products.

The results can be refined by adding further information to the scenario to reflect the impact of external variables (covariates) on the forecast.

Run instantaneous cluster and segmentation analysis

BOARD BEAM can automatically define clusters of similar data objects, by using the k-means methodology.

Users choose the number of groups and the observation sets (e.g. profitability and turnover) and BOARD creates k-means clusters and makes them available as dimensions for analysis in the standard BI and performance management environment.

The seamless integration between the clustering capability and the BOARD environment not only allows for the immediate creation of reports, graphs, and self-service analyses based on clusters, but also enables the use of clusters for forecasting or planning.

Harness the power of BOARD statistical functions

BOARD BEAM offers a set of out-of-the-box statistical functions, from traditional descriptive statistics like min/max, average, standard deviation, to algorithms specifically designed for business analysis such as frequency, recency, dormancy and nascency.

With this seamless integration into the BOARD platform, users can instantly leverage these functions for building any type of analysis, dashboard or report.



3.3 Applications

3.3.1 The toolkit: the fastest way to build analytic applications

Applications (“Capsules” in BOARD terminology) represent the environment where end-users can view, interact with and analyze data. BOARD provides all the tools necessary to build, deploy and maintain these analytic applications. These tools work together within an orchestrated system to provide an intuitive and visual experience for the user. BOARD’s exceptional ease-of-use allows an organization’s non-technical people to become developers, quickly building applications without coding, leveraging all the capabilities offered by the databases and the business logic.

Capsules consist of a collection of screens, with each screen being comprised of one or more BOARD objects. Objects include functionality like: reports, graphs, maps, folders, filters, business rules, gauges and navigation menus. These objects can be freely positioned on the screen through dragging and dropping and then interactively configured by the user.

The objects are automatically synchronized with each other and with the database, allowing users to immediately query, visualize and interact with the application as they build it. The applications leverage BOARD’s powerful business rules engine where users can model any type of business process or calculation (workflow, multidimensional data entry, allocations, consolidations, etc.) using functional business terms and logic.

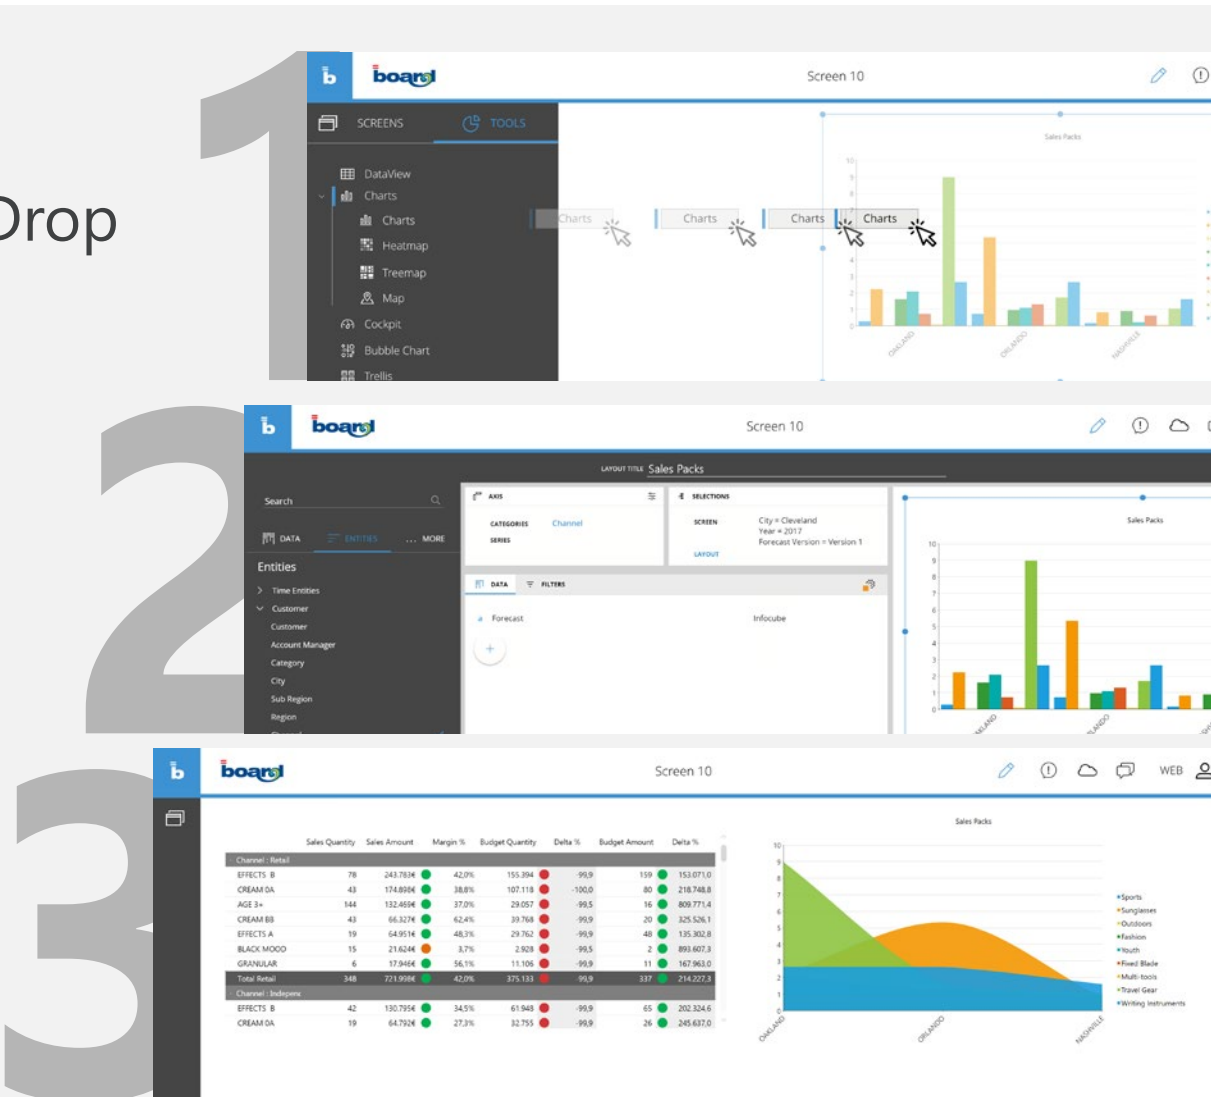
The benefits realized through this approach include:

- Eliminates the need for high levels of technical expertise, giving decision-makers the power to build their own applications
- Enables applications to be perfectly customized to meet the business needs
- Provides flexibility to respond to ever changing demands on the business
- Perfectly fits agile development methodologies
- Substantially reduces the complexity and costs associated with building, customizing and maintaining a comprehensive BI and performance management environment

Drag and Drop

Configure

Run



3.3.2 Storytelling

Users can create personalized versions of application screens (subject to security) and save them in the form of a presentation. Once created, the new presentation can be saved into a personal workspace and shared with individuals and workgroups who can then collaborate on the contents.

Presentations use “live data”. When presentations screens are opened, the data is re-queried and the presentations maintain the same interactivity that they have in the application environment, including drills, pagers, selections, etc. BOARD’s presentation functionality offers a governed but agile and collaborative environment that moves reporting from a pre-canned, rigid and top down discipline into an adaptive and collaborative user-driven process.

3.3.3 In-context Collaboration

User collaboration is greatly facilitated through BOARD's native chat functionality. Users can work together on a shared analysis while interacting in real time in a chat session. The ability to define workgroups completes the picture: virtual teams can instantly discuss their findings, explore new insights and make quick collaborative decisions based on a validated, up-to-date and holistic view of enterprise data.



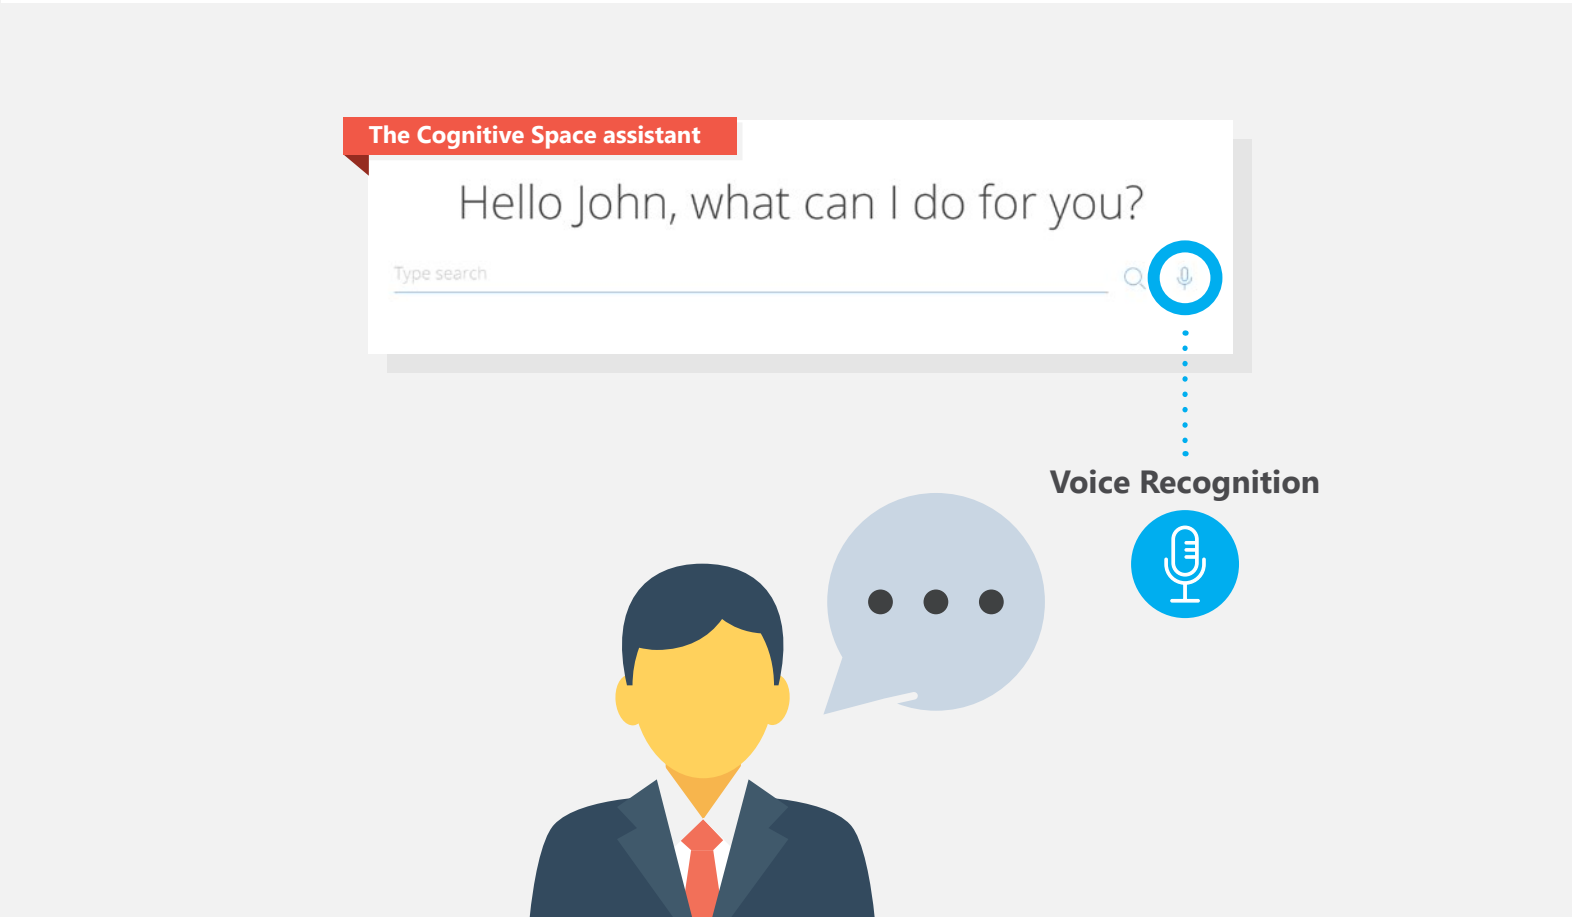
3.3.4 Cognitive technologies

"Cognitive technologies" are more than just fancy words, they represent a collection of technologies that enable people to interact with computers like people interact with each other. With these, come exciting improvements in ease-of-use and user engagement.

BOARD offers the following capabilities, which leverage cognitive technologies, to bring a totally new way to interact with data across the entire BOARD platform. These capabilities are available to all users to the entire platform, cutting across all databases and applications.

Natural Language Recognition (NLR) and Generation (NLG)

To simplify users' access to information, traditional drag-and-drop report creation has been augmented with the ability to directly talk with the system or to interact through search strings, like performing a web search. Behind the scenes, a powerful technology translates speech into words, words into queries, queries into meaningful reports and reports into smart descriptions of the analytical highlights.



Search Generated Analytics

BOARD's Cognitive Search feature makes the creation of reports and analysis as easy as using a web search engine. Users can simply type a question and BOARD, through a fuzzy search mechanism, translates this request into a machine-generated report. BOARD will also evaluate every existing screen in the system and provide a ranked and scored list of screens based how well they answer the question.



Narration

BOARD can generate smart descriptions of your reports, delivered in both text and voice, that highlight the most important and interesting elements. Instead of a business user having to manually review data to spot key highlights and then write commentary, BOARD can automate the entire process, generating smart and engaging content that highlights the key facts represented in the report.

4. CLIENTS



Build once, deploy anywhere

BOARD ensures a consistent experience across web, mobile and desktop environments so developers can build a single application and deploy it to users on variety of devices. BOARD applications can be accessed through the following clients:

- BOARD Web
- BOARD Mobile
- BOARD Office Add-Ins
- BOARD Windows Client

4.1 BOARD Web

BOARD Web leverages HTML 5 technology and was designed for high performance and gesture based interactions. The powerful combination of HTML 5 with the "toolkit" allows users to build applications, without any coding, for immediate on from any device. BOARD will always provide a compelling and consistent user experience, whether accessed from a desktop, a tablet or a smartphone.

Key facts

- Build once without coding, deploy everywhere
- Consistent user-experiences with laptops, tablets and smartphones
- Modern user interface supporting gestures
- Enhanced experience - HTML 5 technology provides more features to web users and supports interactive procedures
- No need for plug-ins

BOARD Web is a zero-footprint, rich-internet client that runs on all popular web browsers, including: Microsoft Edge, Microsoft Internet Explorer, Mozilla Firefox, Apple Safari, Google Chrome, Microsoft Windows and Apple Mac OS X.

BOARD Web allows users to:

- **Access BOARD applications**
- **Edit reports and run interactive analyses**
- **Execute data entry**
- **Build applications (coming in Board 10.2)**
- **System administration and management (coming in Board 10.2)**

Any application developed on the BOARD Web Client can be accessed through either the web or the BOARD Windows Client, making it possible to manage hybrid (web and Windows) deployments without having to build and maintain two separate environments. BOARD Web supports mash-up and can make BOARD inputs available to external applications or acquire inputs from external applications without the need to build integration adaptors. This feature, combined with the Single Sign-On functionality, enables BOARD to be integrated with any company portal.

4.2 BOARD Mobile

BOARD Mobile is a native environment for Windows 8 and iOS based tablets and was designed for high performance and gesture based interactions.

BOARD mobile allows users to:

- **Access BOARD applications**
- **Run interactive analysis**
- **Perform data entry**

Users can also access data in offline mode making it possible to work without a network connection. BOARD Mobile applications are built using the “toolkit” in the same way as standard BOARD capsules and allow organizations to deploy almost any type of application screen to a mobile device in a matter of minutes.

Consistent user experience

BOARD ensures a consistent experience across web, mobile and Windows clients so users can enjoy the same familiar environment on various devices.

Mobile data entry

BOARD Mobile supports online data-entry. Users can input data on their tablet, save and immediately have the changes reflected across all devices. This allows organizations to implement mobile business processes for simulation, planning and forecasting.

Shared insights

BOARD Mobile makes it possible to share reports and visualizations between BOARD Mobile users by sending links. The link recipient will be able to access and modify shared reports and analyses, subject to their security rights. Furthermore, any report or visualization 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. All existing BOARD applications created for a desktop can be easily converted to mobile applications. BOARD includes an automatic compatibility check to reduce development errors, time to solution and deployment issues.

4.3 BOARD Office Add-Ins

BOARD empowers users with self-service analysis and reporting within the MS Office environment. Through the BOARD MS Office add-ins, users can access BOARD multidimensional data directly within Excel, Word and PowerPoint and also leverage traditional OLAP analysis functions, such as drill-down, slice and dice, filtering and ad hoc query.

BOARD objects (e.g. reports, charts and gauges) can be easily embedded within MS 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 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 within the security and data consistency of the BOARD platform.

4.4 BOARD Windows Client

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

BOARD Desktop allows users to:

- Access BOARD applications
- Run interactive analyses
- Perform data entry
- Build applications
- System administration and management

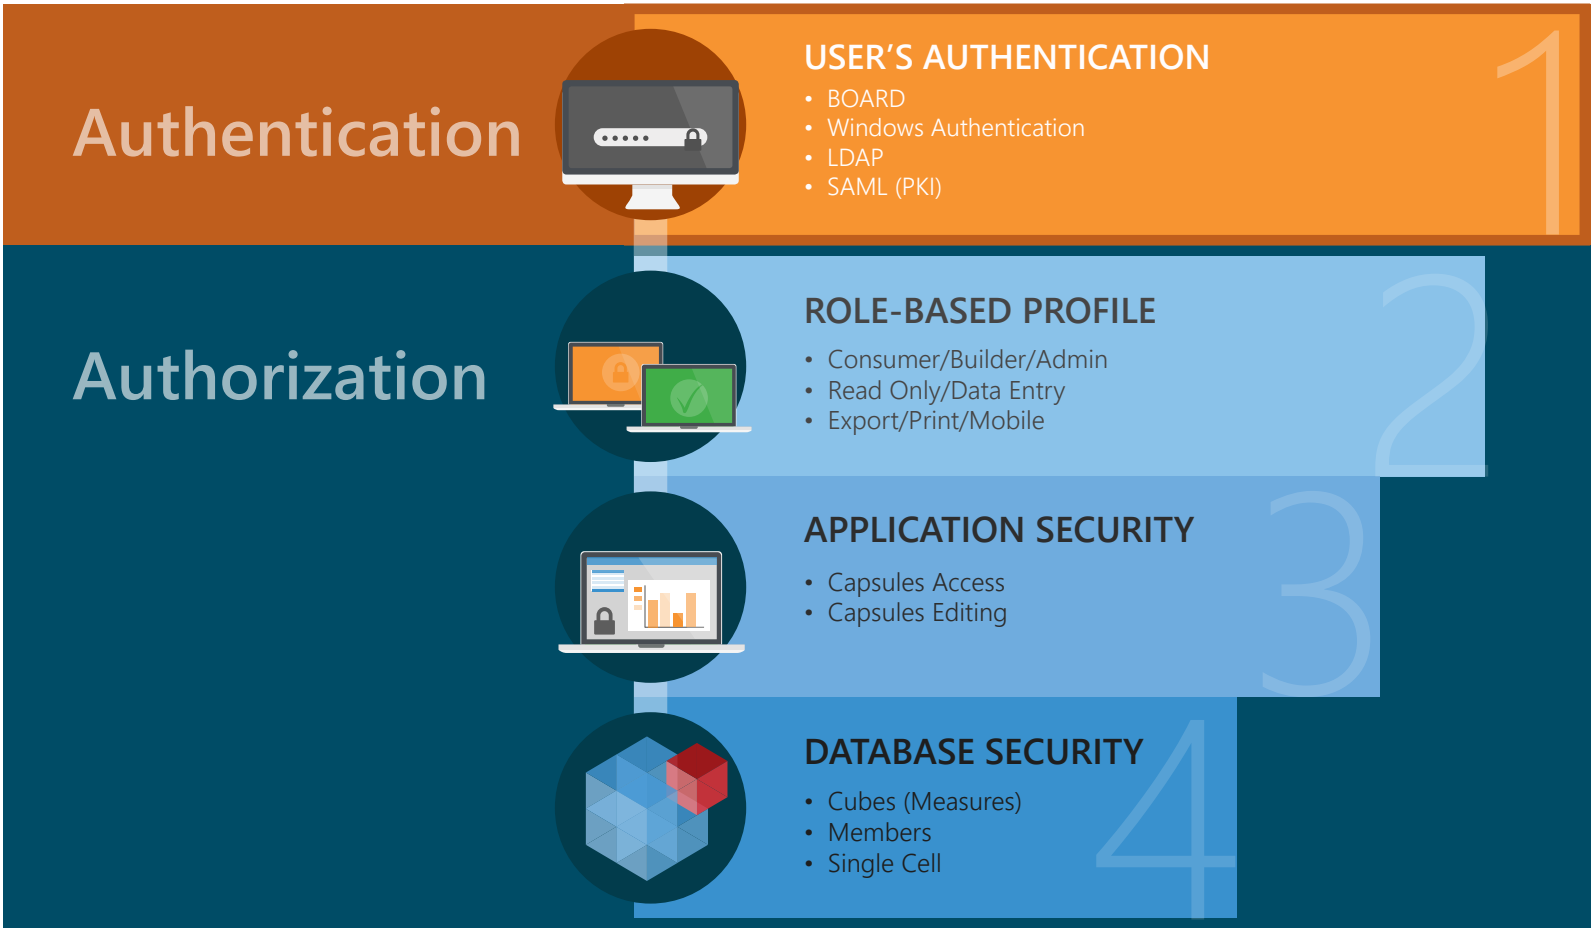
A modern auto update functionality makes it easy to manage BOARD Desktop clients in large scale deployments, eliminating all the technical issues usually associated with Windows client versioning.

5. SECURITY

BOARD's multi-layer security model is built to meet the most critical security requirements of any organization. With BOARD, you can combine data access authorization with application privileges and action permissions to provide pervasive, granular, cell-level security that can be easily personalized for individuals, user roles, and user groups.

5.1 Multi-tier security model

BOARD empowers users with self-service analysis and reporting within the MS Office environment. Through the BOARD MS Office add-ins, users can access BOARD multidimensional data directly within Excel, Word and PowerPoint and also leverage traditional OLAP analysis functions, such as drill-down, slice and dice, filtering and ad hoc query.



5.2 Authentication methods

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

- **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 and maximum retries.
- **LDAP authentication:**
The username and password given by the user is validated against an LDAP server. This method only requires 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.
- **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.
- **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.

5.3 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:

- *Database and Capsule authoring features through the license profile*
- *Restrict or grant access to given Capsule*
- *Restrict or grant access to a BOARD database, specifying a profile for each database*
- *Restrict or grant access to filtering, printing and other*

5.4 Application authorization

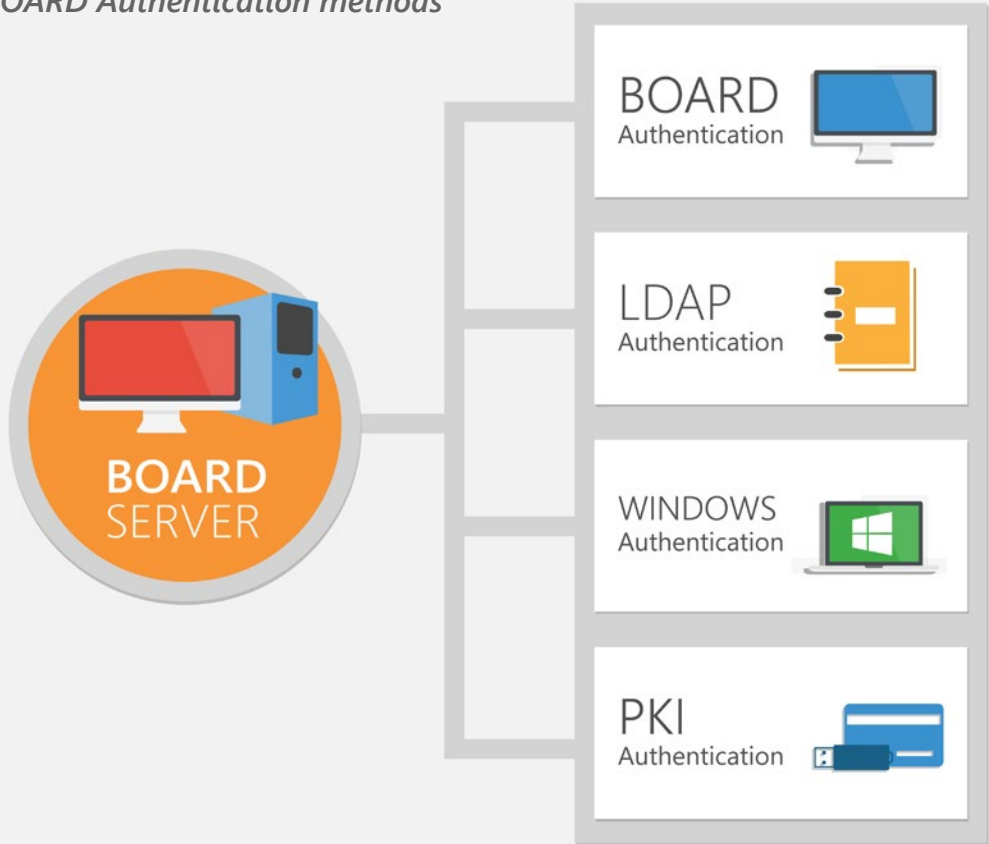
The BOARD Capsules can also embed additional authorizations (e.g. to allow opening or editing a Capsule). This level of authorization ensures that a Capsule can't be opened or altered even if it is copied or moved to a different environment (i.e. a different BOARD server) where the user may have a higher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

5.5 Data access

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

- *Granting/denying access to create or modify a data model: authoring dimensions, cubes, data sources and other database objects*
- *Granting/denying access to cubes with read-only or read-write privileges*
- *Setting filters on dimensions down to cell-level (e.g. to limit a certain security profile to a subset of data using the entities of the data model, such as regions, companies, cost centers or divisions)*

BOARD Authentication methods



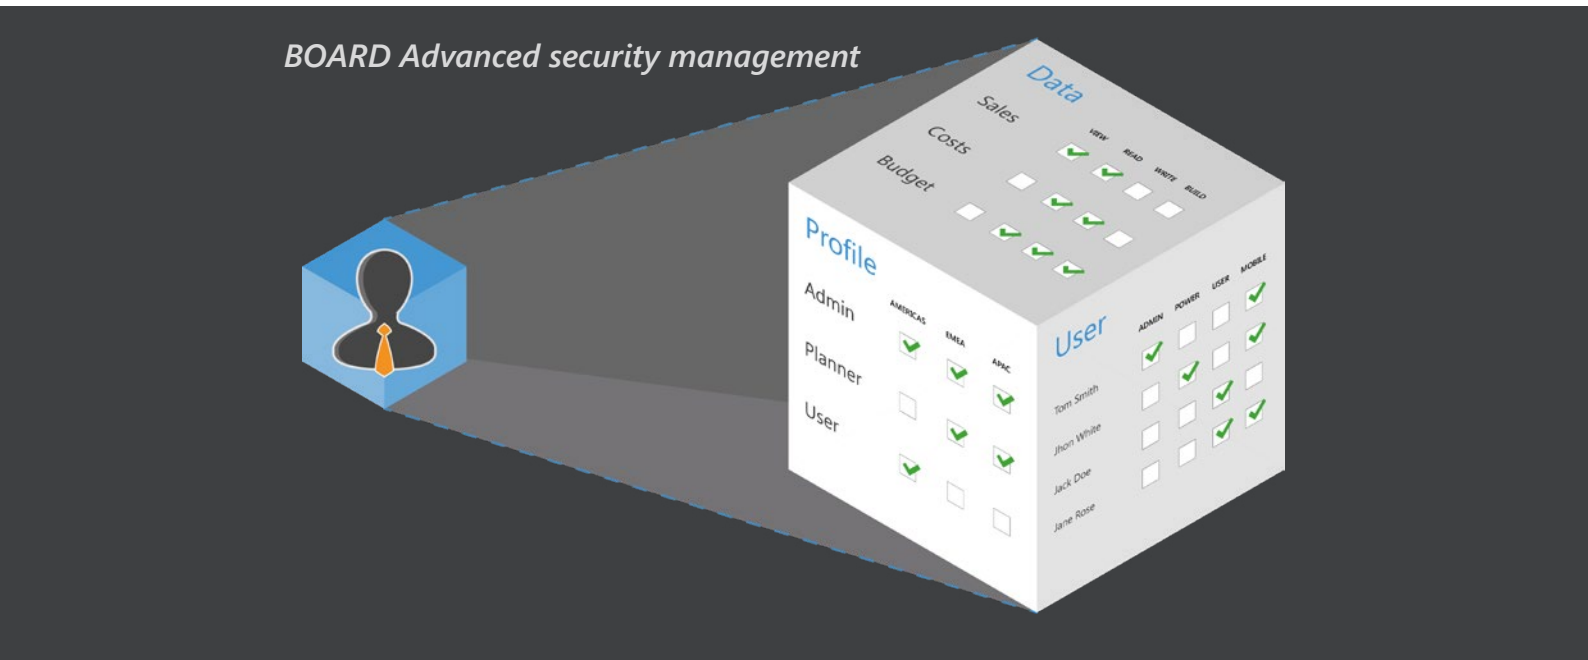
5.6 Advanced security management

BOARD offers the possibility to store and manage authorizations within a BOARD “authorization” database that contains the metadata of all other BOARD databases. A simple idea - use BOARD to manage information (metadata) about BOARD databases and applications and build a completely bespoke solution to handle security 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 within a BOARD application, making it possible to decentralize the security model and to perfectly adapt it to complex organizational structures.

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.



6. INFRASTRUCTURE

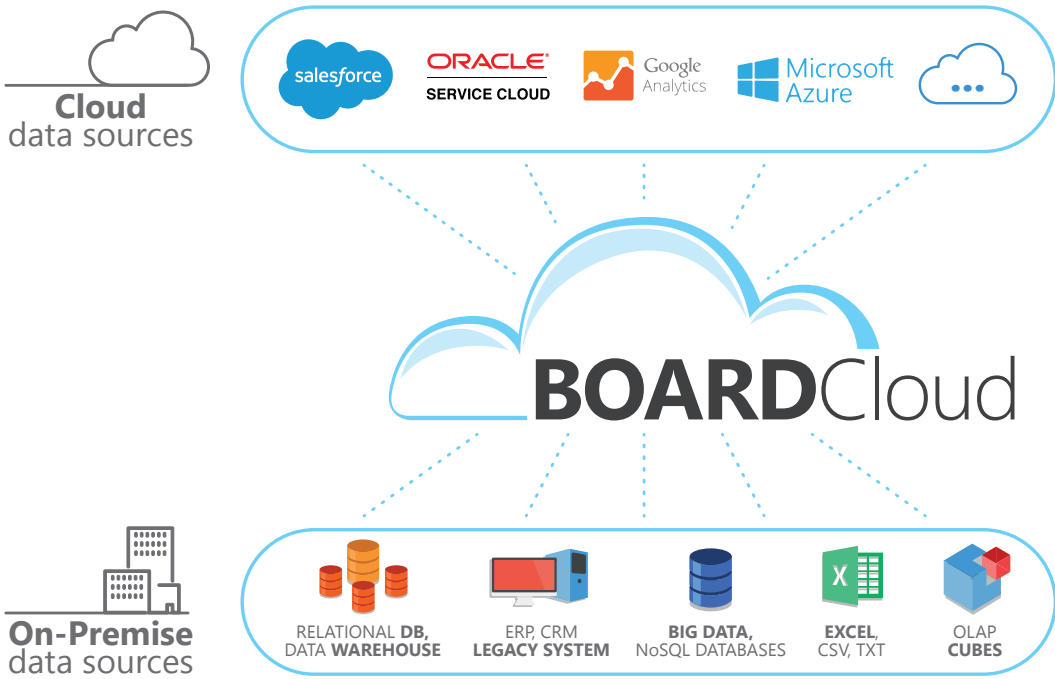
6.1 Cloud or on-premises

On-premises, on-hosting or on-cloud: you can easily implement BOARD with the same reliability and security. With BOARD, you can choose the model that best fits your business at the point of roll-out, having the freedom to move to a different model if and when your needs change.

BOARD Cloud is a SaaS version of the BOARD all-in-one platform. Backed by Microsoft Azure, BOARD Cloud gets your BI and CPM applications up and running faster while providing world-class security, reliability, scalability and performance.

Designed to offer a seamless connection to on-premises and cloud data systems in real time, BOARD Cloud allows you to deliver enterprise analytics, planning and forecasting that fully leverages your existing software investments.

Unlike other solutions on the market, BOARD also supports writing-back from cloud to on-premises, through a secure tunneling connection. This way, your planning data stored in BOARD can easily be integrated back to the underlying legacy systems residing behind your firewall.



Key facts

- **Sign in through a unified portal:** *cloud.board.com*
- **Private BOARD server on Microsoft Azure data centers**
- **Modern, enterprise-grade cloud infrastructure, security and reliability;** *including physical isolation of your data*
- **Hybrid architecture.** *Whether your source data resides in the cloud on on-premises behind your firewall, BOARD Cloud allows you to easily access and integrate them.*
- **Secure connection to on-premises data.** *The BOARD Cloud Connector does not require any firewall changes and immediately creates a secure tunnel between the BOARD Cloud server and your on-premises data.*
- **Write back to relational sources**
- **Native data connectors** *for the most popular cloud data sources*

6.2 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 provides horizontal scalability, supporting both read and write operations. BOARD Cluster provides cutting-edge solutions 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 analyzing very large data volumes that affect 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 BI 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 routine overnight operations such as data feeds and bulk loads from source systems.

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.

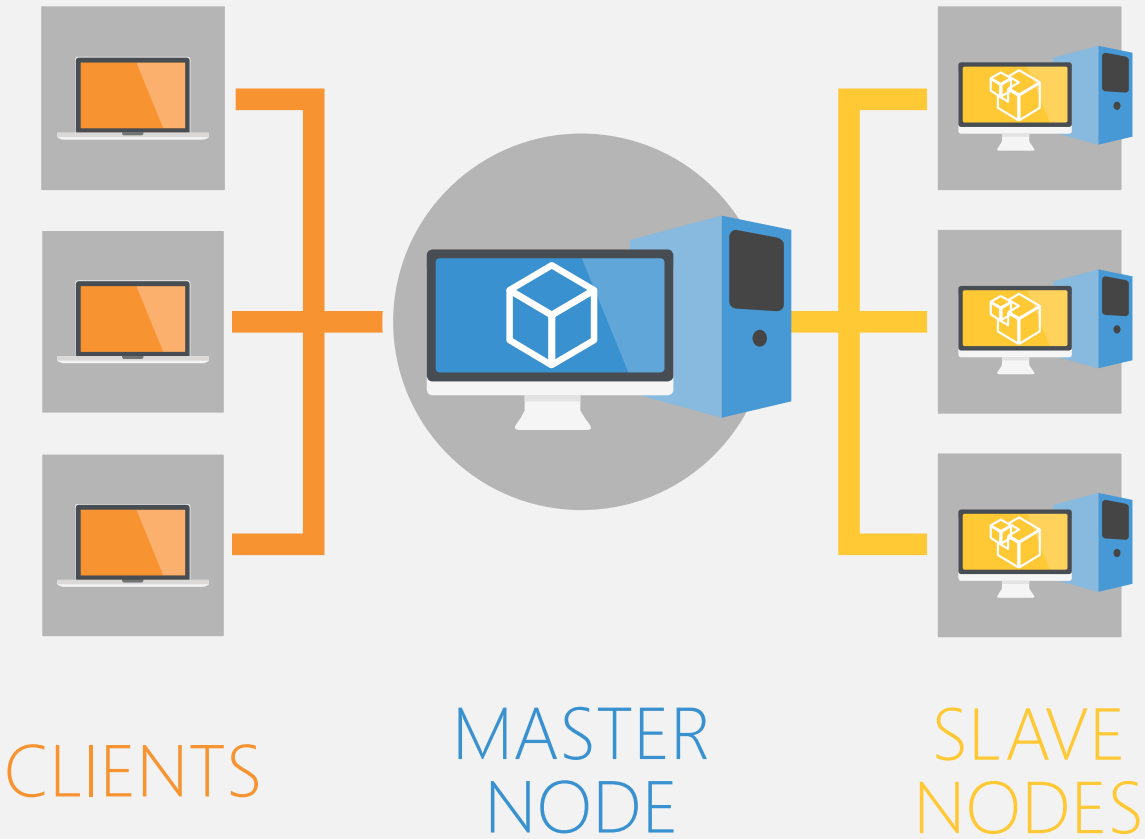
Geographic latency

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 an immediate in-memory update.

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

BOARD In-memory cluster: architectural schema



6.2.1 How the BOARD 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. The slave nodes work locally on their in-memory 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.

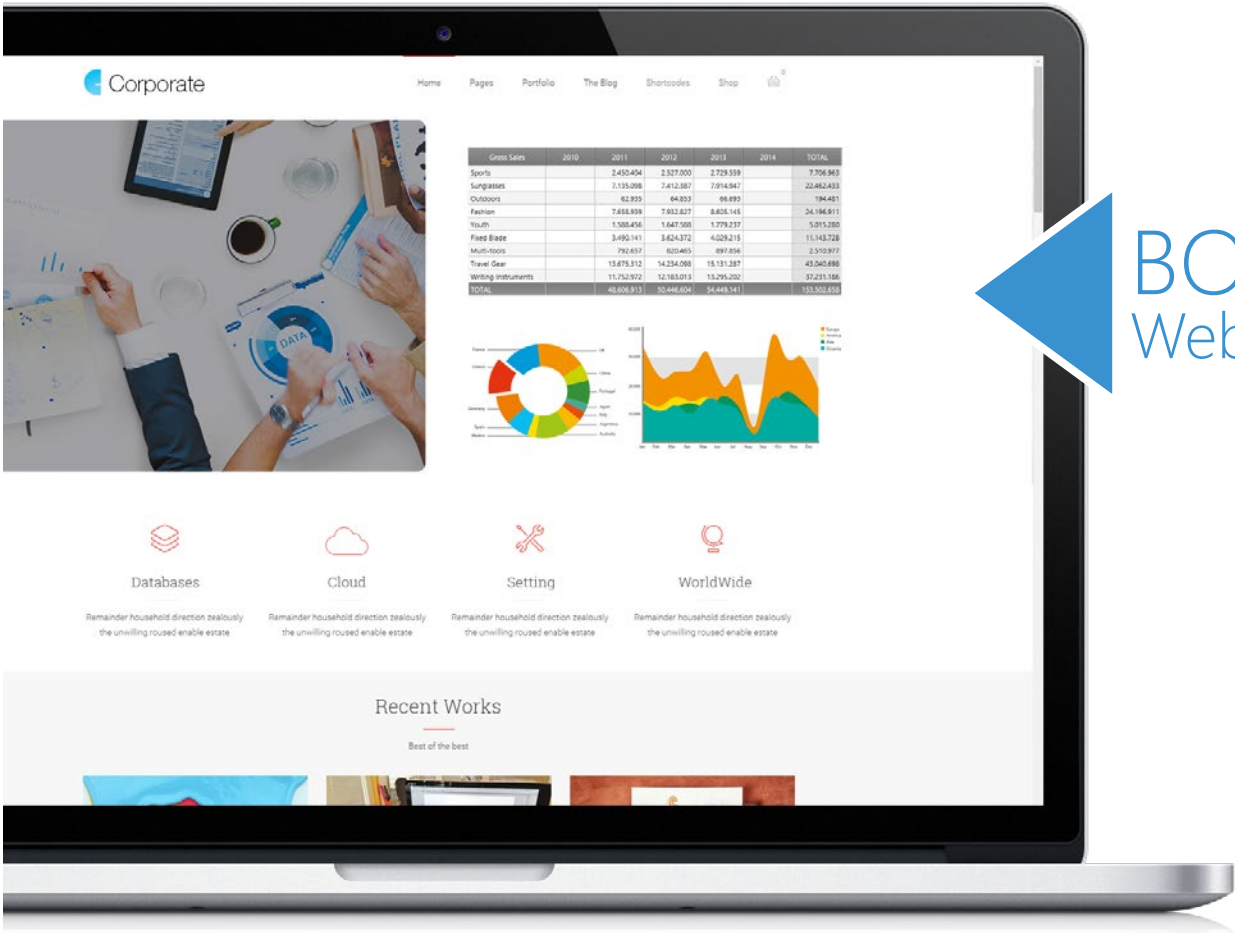
At a given point in time, the data on a slave node can be different (more recent) than the data in another slave node but eventually all the slave nodes become consistent (Eventually Consistent Model). The whole process is performed in RAM using the BOARD HBMP in-memory technology, thus making it extremely 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 synchronized between servers through a rapid in-memory update.

6.2.2 Adaptive load balancing

BOARD Cluster can be configured in several ways to adapt to different use cases:

- **Users:** BOARD Master Server automatically redirects users to the various slave nodes. This approach is suited to handle a large number of users when there are no other specific architectural requirements.
- **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).
- **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' workloads based on the performance of the underlying hardware architecture.
- **Ad-hoc configuration (Specialized Nodes):** When managing geographically distributed, complex planning and simulation processes, an automatic load-balancing configuration is often an ineffective architectural choice. The system administrator needs the ability to define which data, users and processes to handle on any single node in order to align the system architecture with the specific functional requirements, resolve potential workload bottlenecks and optimize the performance of the entire process.

6.3 Interoperability and portal integration



BOARD
Web Mash-up

6.3.1 Web mash-up and portal integration

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

Through Java Script it is possible to call the methods exposed by the BOARD Web Client and interact with BOARD applications. These capabilities combined with the single sign-on functionality enables BOARD to be integrated with any company portal.

6.3.2 SharePoint integration

BOARD Web and SharePoint are both based on a Service Oriented Architecture (SOA) and Microsoft .Net Framework. This technology alignment makes the integration between the two systems quick and easy.

The integration can be handled at three different levels:

Web page viewer

The simplest option to render 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 approach 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 longer to return the content than the other parts.

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”).

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 BOARD Web Client. 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 that point.

This solution permits a comprehensive integration between BOARD and the SharePoint environment but will require the use of HTML and JavaScript programming.

6.3.3 BOARD web services

BOARD Web Services allow flexible integration of BOARD with external applications. BOARD was designed to be used as a web service and manages SOAP messages (usually formatted in XML) and returns XML data structures.

The services supplied by BOARD can be used by any application in a simple and standardized way (W3C Standard), including: other web services, legacy systems, Java, .NET applications, web applications and by company portals.

By providing these services, BOARD can extend advanced BI and performance management functions to the existing company application environment. BOARD Web Services are based on two types of methods:

- Data dictionary functions: (database metadata browser) that provide the list of Cubes (Measures), Dimensions, Hierarchies and Members
- Query execution functions: allowing access to data

Where to find BOARD

BOARD INTERNATIONAL Headquarters

EUROPE / APAC	AMERICA
Piazza Bernasconi 5, 6830 Chiasso, Switzerland	33 Broad Street, Suite 200 Boston, MA 02109
Tel: +41 916975410	Tel: +1 6172032173
Fax: +41 916975429	Fax: +1 7812070347
e-mail: info@board.com	e-mail: infoamericas@board.com

ARGENTINA	ASIA	AUSTRALIA
BOARD Argentina	BOARD Asia Pacific	BOARD Australia
Maturin 2870	60 Paya Lebar Road	100 Walker Street
Ciudad de Buenos Aires (C1416CTP) - Argentina	Paya Lebar Square, #08-23 Singapore 409051	NSW 2060 North Sydney
Tel: +5491164193822	Tel: +65 62789138	Tel: +61 289049777
e-mail: latam@board.com	Fax: +65 63851690	e-mail: australia@board.com
	e-mail: asiapacific@board.com	
BENELUX	FRANCE	DACH
BOARD Benelux	BOARD France	BOARD Deutschland
Karspeldreef 14	21 Rue Auber	Schaberweg 28
1101 CK Amsterdam	75009 Paris	61348 Bad Homburg v.d.H.
The Netherlands	Tel: +33 143873736	Tel: +49 6172171170
Tel.: +31 20 7630591	e-mail: infofr@board.com	Fax: +49 61721711770
e-mail: infobenelux@board.com		e-mail: infode@board.com
HONG KONG - CHINA	INDIA	ITALY
BOARD Hong Kong	BOARD India	BOARD Italia
Unit A, 20/F	602, Sixth Floor,	Via Caduti di Marcinelle 5
9 Queen's Road Central, Central	Landmark Pllatinum	20134 Milan
香港皇后大道中9號20樓A室	Plot No.337/338	Tel: +39 022108071
Tel: +852 31897087	Khar West, Mumbai 400052	Fax: +39 0221098550
Fax: +852 31897689	Tel: +91 22 26050100	e-mail: infoit@board.com
e-mail: hongkong@board.com	e-mail: india@board.com	
JAPAN	MEXICO	MIDDLE EAST
BOARD Software K.K.	BOARD Mexico	BOARD Middle East
Level 27 Shiroyama Trust Tower	Durango 269, Piso 3.	Level 29, Marina Plaza
4-3-1 Toranomom Minato-ku	06700 Col. Roma Norte	Dubai Marina,Dubai,UAE
Tokyo 105-6027	Tel: +52 5585968693	P.O Box 112229
Tel +81 03-5403-4660	Fax: +52 5552720178	Tel: +971 4 559 5651
Fax +81 03-5403-4661	e-mail: info@board.com.mx	e-mail: infome@board.com
e-mail: japan@board.com		
SPAIN - PORTUGAL	UK - IRELAND	USA - CANADA
BOARD Ibérica	BOARD UK	BOARD Americas
Avda. Portal de l'Àngel, 36	Tower 42, Floor 20	33 Broad Street, Suite 200
08002 Barcelona	25 Old Broad Street	Boston, MA 02109
Tel: +34 93 4925702	London, EC2N 1HQ	Tel: +1 6172032173
e-mail: infoes@board.com	Tel: +44 2037272222	Fax: +1 7812070347
	Fax: +44 1908255891	e-mail: infousa@board.com
	e-mail: infouk@board.com	

REQUEST A DEMO!

www.board.com

AMSTERDAM, AUSTIN, BARCELONA, BOSTON, BUENOS AIRES, CHIASSO, DUBAI,
FRANKFURT, LIMA, LONDON, MADRID, MEXICO CITY, MILAN, MUMBAI,
MUNICH, PARIS, SINGAPORE, SYDNEY, TOKYO