



The PivotLink Solution Framework

Delivering self-service business intelligence to decision-makers

“PivotLink makes it possible for people without technical expertise to explore and experiment with their data in order to answer questions they have about their business.”

EXECUTIVE SUMMARY

Today’s workforce expects business information to be as accessible and easy-to-use as Google or online travel. But behind the scenes, IT managers need to wrestle with massive amounts of data from inside and outside the organization to try to make this information available. Even with the considerable advances in computing during the past 15 years, business intelligence (BI) implementations continue to be the most challenging IT projects, wrought with complexity, low user adoption and poor ROI.

Current BI practices typically result in a proliferation of complex information silos that are hard to manage and rarely help create an enterprise-wide view of the business. Worse yet, traditional BI solutions often become the source of information bottlenecks and the dreaded “multiple-version of the truth” problem that they are intended to fix.

Companies need affordable, easy-to-deploy, and easy-to-use analytic tools that can quickly deliver real value to a wide range of knowledge workers. PivotLink’s approach to BI is a paradigm-shifting model that combines the advanced technologies of in-memory analytics, columnar data storage, cloud computing, and Software-as-a-Service (SaaS) delivery to provide a more dynamic analytical capability for decision-makers. The solution is designed to align with the way people work and consume analytical information.

This white paper describes the PivotLink Solution Framework and how it delivers self-service business intelligence to drive analytical competence and gain better business insight without IT assistance.

THE CHALLENGE: INTEGRATING DATA IN AN ANALYTICAL FRAMEWORK

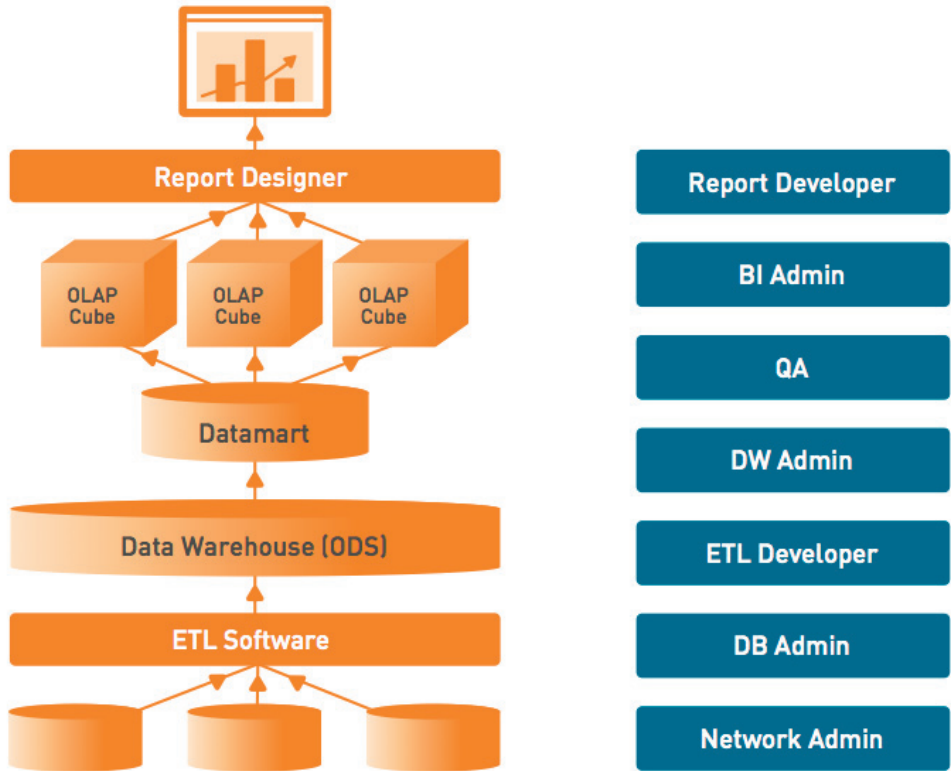
It's extremely challenging for organizations to integrate data from a variety of systems and manipulate it to make decisions that shape business strategy. Today's business users require nearly instantaneous access to information from inside and outside the organization, and the ability to turn this information into actionable intelligence.

The need for context-rich information is at the core of business intelligence (BI). The past three decades have seen the development of several BI tools to address this challenge (Figure 1). However, these tools were rooted in the technology and processes of operational and transactional data management systems. They didn't provide an analytical framework for context-based performance indicators.

With these limitations, traditional BI systems couldn't efficiently deal with the pace of modern business. New data sources and the ever-increasing volumes of data couldn't be added quickly to the BI infrastructure. Data from incompatible sources couldn't be integrated easily. New report requirements didn't fit into the fixed models of the BI systems. Decision makers needed more reliable, consistent, and timely information at all levels of granularity.

It's important for businesses to aggregate data from different transactional sources in order to create reports and measures such as key performance indicators (KPIs) for gauging organizational performance. These reports tend to be highly complex, presenting data across several dimensions (for example, sales revenue by region). The challenge is to provide this information dynamically, in a context-rich analytical framework, so that business users without technical expertise can explore and experiment with data in order to answer questions they have about their business.

Figure 1: The traditional business intelligence solution model



THE PAIN: UNLOCKING DATA FROM TRADITIONAL MODELS

The key to unlocking data for exploration and experimentation is to release it from restrictive, predefined relationships and management systems. At the core of a traditional BI system is the underlying schema that stores data in tables and fields and defines data relationships. These generally-accepted models are known as entity-centric, star, and snowflake. Derived from transactional processing systems in which data integrity and moment-in-time accuracy are essential, these models enforce a structured analysis scenario.

While there are advantages to using these models (Table 1), they are only able to support scenarios in which the reporting needs are straightforward, such as snapshots of particular metrics – without the need for “slicing and dicing” the data.

More complicated reporting needs, such as changes in several metrics over time, require cost-prohibitive implementation efforts.

Traditional BI tools are also limited by their underlying data management capabilities. Their data processing models (Table 2) are also a legacy of transactional processing systems. Data management enables efficient (physical) data storage, organization and retrieval, but traditional BI tools don’t develop these to appropriately support analytical processing needs.

To go beyond these limitations, it takes skilled and dedicated information technology professionals sitting with staff from the business units and launching several data warehousing and reporting projects. Despite significant advances in hardware and computing, efficient data processing and retrieval for analytics continues to be a challenge.

Table 1: Advantages and disadvantages of traditional data models

| ADVANTAGES | DISADVANTAGES |
|--|---|
| <p>They enforce moment-in-time accuracy of data relationships</p> <p>They can enforce business rules with referential integrity</p> <p>They provide a simple mechanism to manage low-complexity reporting requirements</p> | <p>They fail to adapt to changing data relationships over time</p> <p>They don’t guarantee aggregation integrity when referential integrity can’t be enforced from a historical perspective</p> <p>They freeze reporting requirements into the fixed dimensions of the model, making change expensive</p> <p>Modeling is time- and people-intensive</p> <p>They can’t support complex reporting needs if they make assumptions about relationships that don’t exist, such as artificial constructs to support the model</p> |

Table 2: Limitations of current data processing models

| RDBMS | CUBES | HYBRID MODEL |
|--|--|--|
| Row-oriented storage impedes efficient access of data elements for analysis. | Pre-designed hierarchies prevent flexible navigation paths to information. | There are too many points of redundant work. |
| Requirements in hardware and storage increase exponentially to accommodate different users and changing data volumes, complexity, and sources. | Extensive data normalizations and other transformations make it harder for business users to understand the results. | The model makes it even more difficult to enforce a single version of the truth. |
| Query processing engine is designed to optimize transaction handling, not to optimize aggregation efficiency and integrity. | Extensive use of predefined data and excessive indexing makes the solution vulnerable to failure and lengthy delays in availability. | The model promotes silos of information instead of an integrated view. |

THE PIVOTLINK SOLUTION FRAMEWORK

Addressing these BI challenges in today's information intensive organizations requires a fundamentally different perspective on data modeling and processing. The PivotLink Solution Framework addresses these challenges with the following components (numbered in Figure 2):

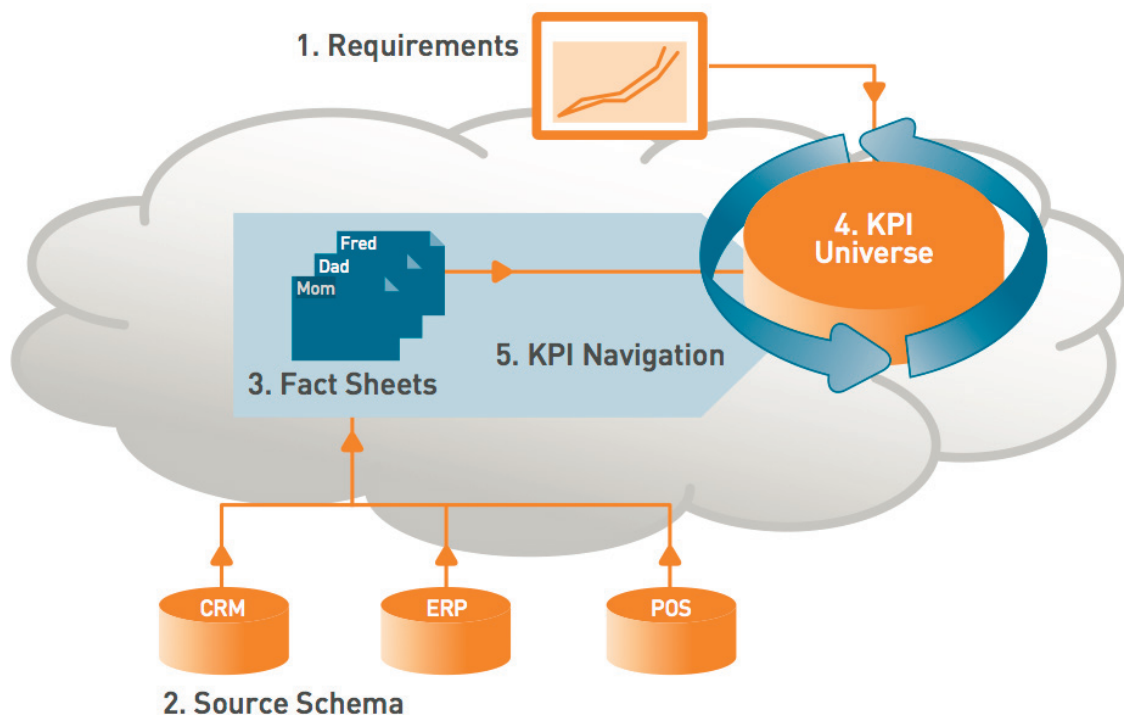
1. **Requirements:** the analytical requirements for key performance indicators (KPIs) supporting business decision making.
2. **Source schema:** customer data in native format from multiple (or single) sources.
3. **Fact sheets:** the source of underlying calculations and data supporting metrics and KPIs.
4. **KPI Universe:** a set of metrics and analysis components.
5. **KPI Navigation:** a mechanism that enables a multi-dimensional perspective on data.

Schema-independent modeling

Unlike traditional BI approaches, the PivotLink Solution Framework is independent of any schema. The storage of data in fields and tables, and their associated relationships, are based on the "natural" state of the data. Artificial key relationships and data integrity rules based on restrictive and volatile hierarchical relationships are not enforced. This allows data to be viewed in an adaptive fashion, instead of the snapshot views supported by the other schemas.

In addition, data modeling is driven from the perspective of facts. (This concept is utilized in traditional BI, but with a different outcome in the PivotLink Solution Framework.) Facts represent the logical definitive source for computing one or more primary metrics. Primary metrics (and composite metrics that are combinations of primary metrics) are the logical concepts modeled in the PivotLink framework. These concepts ensure that there is a single version of the truth and the highest level of aggregation integrity when analytics are produced, without the artificial constraints introduced by traditional BI tools.

Figure 2: The PivotLink Solution Framework

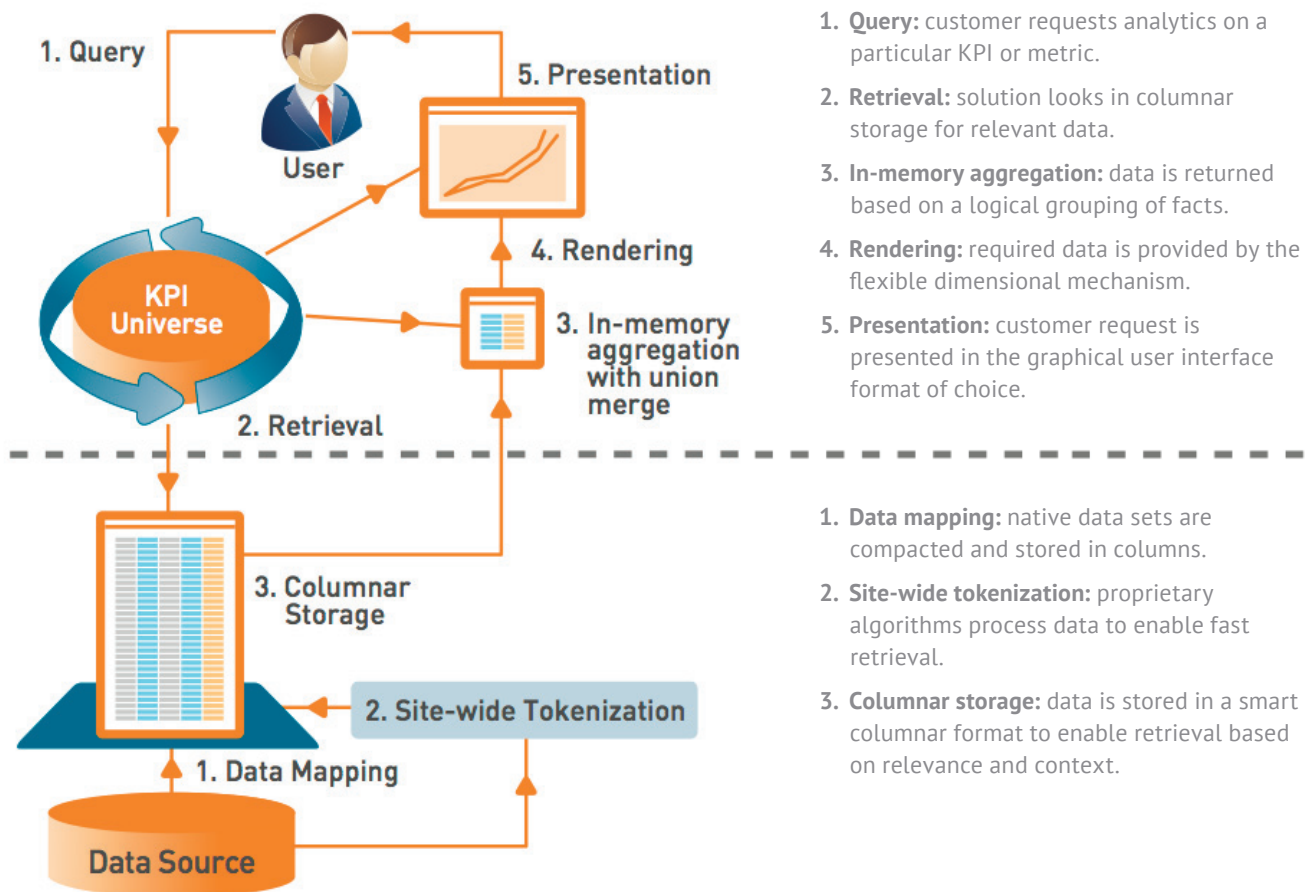


The framework provides an efficient mechanism for analytics. The specific join relationships are not spread across tables outside the context of analytical measures. As a result, a business user can view data over time, and with any level of granularity, without requiring development work or changes to data output from the transactional systems. The framework offers greater elasticity in exploring and analyzing information, with complete control over this iterative process. Enterprises gain the benefits of business intelligence without any disruption to their core IT environments or processes. The PivotLink Solution Framework delivers on the promise of business intelligence by converting data into actionable information.

Smart data processing

The PivotLink solution’s smart data processing focuses on relevant data, rather than all data (Figure 3). Proprietary technology enables compact data storage and retrieval from numerous source systems – in direct contrast to legacy BI systems. The PivotLink solution uses site-wide tokenization to process large volumes of data for quick retrieval. Data sets are compacted and stored in columnar format. Disparate data sources can be accommodated and efficiently merged to provide quick, flexible analytical results on-demand.

Figure 3: PivotLink smart data processing



- 1. **Query:** customer requests analytics on a particular KPI or metric.
 - 2. **Retrieval:** solution looks in columnar storage for relevant data.
 - 3. **In-memory aggregation:** data is returned based on a logical grouping of facts.
 - 4. **Rendering:** required data is provided by the flexible dimensional mechanism.
 - 5. **Presentation:** customer request is presented in the graphical user interface format of choice.
-
- 1. **Data mapping:** native data sets are compacted and stored in columns.
 - 2. **Site-wide tokenization:** proprietary algorithms process data to enable fast retrieval.
 - 3. **Columnar storage:** data is stored in a smart columnar format to enable retrieval based on relevance and context.

By enabling the viewing of data outside the context of transactional and operational processing systems, the PivotLink Solution Framework aligns with the way people work and consume analytical information.

PIVOTLINK: GAINING INSIGHT ON DEMAND

Organizations face the challenge of providing analytic tools and access to data to a wider range of business users in a way that is highly relevant for decision-making. Traditional BI tools are very limited for supporting true business analytics, due in part to their origins in transactional and operational processing systems.

The PivotLink solution, based on new technologies, can fully address the needs of today's information-intensive organizations. Its paradigm-changing model provides business analytics as a service rather than as an on-premise cost-center. This model solves the cost, risk, and adoption issues associated with legacy BI systems. The process of developing, deploying, and maintaining analytic solutions is sourced to the Internet "cloud" reliably and securely, resulting in faster time to value, increased scalability, and significantly lower cost and resource requirements. The Software-as-a-Service (SaaS) delivery model unburdens IT departments and empowers business users by enabling them to access their performance measures in meaningful, easy to customize formats over the Internet.

The PivotLink approach eliminates the need for complex data transformations, thereby improving the time to value and the total cost of ownership as compared to other BI approaches. A much broader range of businesses can take advantage of the following benefits of the PivotLink solution:

- Reduce the cost and complexity of BI deployments
- Unburden IT from time-consuming BI projects and shift its focus to data quality, service, and innovation
- Extend more data to more users in a self-service model
- Enable secure collaboration within and outside the organization

PivotLink's self-service BI solutions are designed for the way people work and consume analytical information in today's dynamic and information intensive business environment. PivotLink's dashboards, metrics and analytics make it easier than ever for business users to tap into data online, get better insight into day-to-day activities, make decisions confidently and react quickly to changes in the business. PivotLink is the only proven BI solution delivered as a service with:

- 15,000 users, with 1.3 million ad-hoc reports served every month
- Enterprise class deployments with three billion rows of data
- Customers with multiple data sources
- Pervasive deployment with thousands of users inside and outside the enterprise

PivotLink was founded on the simple, yet powerful principle that it should be easy for business users to securely analyze any data, any way they want, and share their insights with colleagues and partners wherever they are. PivotLink's approach is to put affordable, secure and easy-to-use analytic tools into the hands of the line of business users and empower them to make better, timelier decisions – without burdening IT.

LEARN MORE

Learn how PivotLink's on-demand business analytics can improve your business by visiting our website at pivotlink.com.

ABOUT PIVOTLINK

PivotLink, headquartered in San Francisco, CA, is the leader in on-demand business intelligence (BI). Built from the ground up by Software-as-a-Service (SaaS) and BI industry visionaries and veterans, PivotLink's award-winning solutions put affordable, secure and easy-to-use analytic tools into the hands of line of business users, freeing IT to focus on strategic business initiatives.

PivotLink Corp. | One Post St., Suite 900 | San Francisco, CA 94104
Phone: 1.866.625.9884 | Fax: 415.544.5295 | pivotlink.com