

Whitepaper

4 Steps to Successfully Evaluating Business Analytics Software

Introduction

The goal of Business Analytics and Intelligence software is to help businesses access, analyze and visualize data, and then communicate those insights in meaningful dashboards and metrics. Unfortunately, the reality is that the majority of software options on the market today provide only a subset of that functionality. And those that provide a more comprehensive solution, tend to then lack the features that make it user-friendly.

With a crowded marketplace, businesses need to go through a complex evaluation process and make some fundamental technology decisions before selecting a vendor. Finding a business intelligence (BI) software that will scale with your organization's needs may seem like an impossible task. Here are the four questions you can ask when beginning the BI evaluation process that will save you a lot of time and help set you in the right direction.

I. How Do I Select A Software Stack?

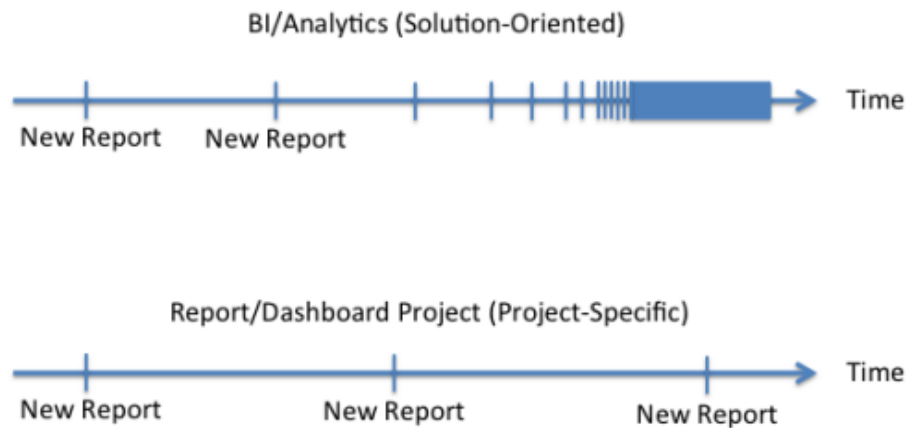
To answer this question, first you need to better understand the options for extracting intelligence from your data. It is important to differentiate between *business analytics* and *dashboard reporting*.

Business intelligence (or business analytics) is a set of theories, methodologies, processes, architectures and technologies that transform raw data into meaningful and useful information for business purposes.

By contrast, dashboard reporting projects have a more limited scope and generally address current requirements rather than future ones. Dashboards and reports tend to be static, created once and simply refreshed with updated data, as opposed to business analytics, which allows an organization to create new reports and dashboards as required.

Project-Specific vs. Solution-Oriented Approaches

Organizations that choose a project-specific approach generally require a short-term solution or have projects of limited scope. Reporting requirements are predictable and consistent, and can be handled with static reports based on single or simple data sources.



Organizations that have a broader view of their requirements will use a longer-term, solution-oriented approach. These organizations understand that their need for reporting will get more complex as data accumulates and gets more complex, and as the rate at which new reports are introduced will increase over time due to integrations, customizations, time for training and increased adoption.

Businesses that opt to use a solution-oriented approach need software that takes into consideration current *and* future reporting requirements. The most effective way to handle this in terms of technical architecture and scalability is by implementing a data warehouse and possibly data marts. (See below).

Technical Overview

If you take a forward-looking approach and see that your business analytics requirements will develop over time, implementing a data

warehouse architecture will allow your solution to scale with your needs. That being the case, your business analytics architecture will probably look similar to the diagram on the right. You can see all data from all data sources is pulled into a centralized database through a process called Extract-Transform-Load (ETL). This database is called a data warehouse (DW).

A data warehouse should be viewed as your company's 'single source of truth' by containing a compilation of only clean and accurate data.

Data warehouses can house historical data as well as current data.



Take note: Management of a data warehouse is often handled by the IT department, limiting the access of all other business groups within the company. Instead, individual teams or business groups may access data in the data warehouse via data marts.

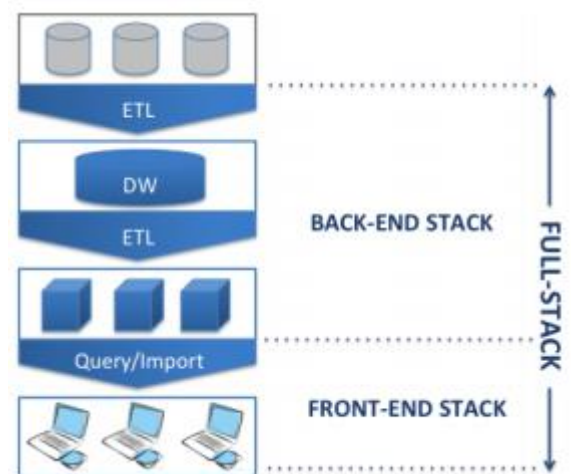
A data mart is a subset of the data warehouse that is oriented to a specific business group or team. Whereas data warehouses have an enterprise-wide depth, the information in data marts usually is limited and only pertains to a single department.

Both data warehouses and data marts refer to the back-end of the business analytics software stack. This is where the data is stored, transformed and managed. Neither the data stored in a data warehouse nor that in a data mart is accessible by end users. In order to extract the data from the back-end, you will also need a front-end visualization tool for data discovery and visual analysis.

Different Types of Software Stacks

Business analytics software is offered in three different configurations.

- ▶ **Back-End Software Stacks:**
Provide only back-end functionality such as data storage, transformation and management (i.e., data warehouse and data mart functionality, as well as ETL capabilities).
- ▶ **Front-End Software Stacks:**
Provide only front-end (end-user facing) functionality, such as data visualization and visual analysis.
- ▶ **Full Software Stacks:** Deliver both back-end and front-end functionality.



Full Stacks vs. Front-End Stacks

If you have not gone through the process of collecting, centralizing and standardizing your data, using a front-end stack instead of a full stack has several negative implications:

Front-End Stack	Full Stack
Detailed Data	Detailed Data
Dirty Data	Clean Data
Non-Centralized	Centralized
Data Mart Capacity or Excel/CSV	Data Warehouse Capacity

The advantages of using a full stack as opposed to a strictly front-end stack are:

- ▶ Enables centralized data storage and management
- ▶ Allows for consolidation of multiple data sources
- ▶ Scales to far larger data volumes and more concurrent users
- ▶ Minimizes the number of data marts to develop and manage
- ▶ Avoids need to create and work with multiple (and potentially huge) CSV extracts
- ▶ Protects operational systems from damage (rogue queries)
- ▶ Allows access to data sources that cannot be queried

When Would You Need the Full Software Stack?

In most cases, implementing only a back-end stack or a front-end stack would not suffice to ensure a real, effective and scalable business analytics solution. In particular, you would need a full stack when:

- ▶ End users want to access centralized data and maintain a single version of the truth rather than build a solution around CSV/Excel extracts.
- ▶ ETL functionality is required, which often happens when multiple data sources are involved and when the data is especially dirty or especially large.
- ▶ The data sources cannot be directly queried, either because they are not supported or because they are part of a critical operational system.

2. How Do I Select The Right Database Technology?

Your next step is selecting the technology configuration that suits your current business needs and will scale for future requirements.

Are you going to use a data warehouse scale ('big scale') technology, or will data mart scale technology be sufficient? To determine the database technology for your business analytics project, you will need to consider: the volume of your data and the rate of increase, the number of users now and in the future, the number of data sources, the complexity of the data, if there's a need for ETL, and the scale and scope of the project.

Data Warehouse vs. Data Mart Technology

Database technologies are used as the backbone for data warehouses as well as for data marts. Since data warehouses typically store far more data than data marts, data warehouse technologies focus on scale.

On a single server, data warehouses are expected to scale to terabytes of data (1 terabyte = 1,000 gigabytes), while data mart technologies only work well in the gigabyte range.

Other points of comparison are listed below:

	DW Technology	DM Technology
Raw data supported (1 mart)	Terabytes-Petabytes	Gigabytes
Number of fields (1 mart)	Practically unlimited	Limited
Number of records (1 table)	Billions	Millions

These differences also demonstrate why data warehouse technologies are often used for data marts, especially when the data marts are expected to grow in size. The inverse is rarely true because of scale limitations. Using data mart technology for a data warehouse scale project generally won't work as a long-term solution.

Data Warehouse Technologies

There are three types of database technologies used for data warehousing:

Software Appliances:

A software application that might be combined with just enough operating system (JeOS) for it to run optimally on industry standard hardware (typically a server) or in a virtual machine. Examples of Software Appliances are: Microsoft SQL Server, MySQL and Sisense.

Computer (Hardware) Appliances:

A computer appliance is a separate and discrete hardware device with integrated software (firmware) designed to provide a specific computing resource. Computer appliances are not designed to allow the customers to change the software, or to flexibly reconfigure the hardware. Examples of hardware appliances are: IBM Netezza and Oracle ExaData.

Distributed Databases:

A distributed database may be stored in multiple computers, located in the same physical location, or may be dispersed over a network of interconnected computers. A distributed database system consists of loosely coupled sites that share no physical components (such as disk, RAM and CPU). Examples of distributed databases are: EMC GreenPlum, HP Vertica and Hadoop.

The summary of differences between the different types of data warehouse technologies is listed below:

	Commodity Hardware	Proprietary Hardware	Distributed Databases
Hardware Class	Commodity	Proprietary	Commodity
Best Architecture	1 server	1 server	Unlimited servers
Capacity	Terabytes	Terabytes	Petabytes
Hardware Costs	4-5 figures	6-7 figures	5-6 figures

Data Mart Technologies

Two primary data mart technologies are used today:

OLAP (Online Analytical Processing):

An OLAP Cube's main purpose is to enable fast performing multi-dimensional slicing and dicing. OLAP achieves fast performance by pre-calculating metrics (field aggregations) for all sets and subsets of unique values in all dimensions (fields) 'over-night'. This avoids performing these slow operations in real-time during the workday. Storing the results of these pre-calculations takes exponentially more storage resources than the actual raw data does, limiting the size of raw data that can make up a cube to gigabyte scale.

IMDB (In-Memory Databases):

Similar to OLAP, IMDB's primary purpose is fast performance. It achieves fast performance by loading the entire data mart into RAM, thus avoiding slow disk-reads ("I/O Bottlenecks"). The size of data mart is effectively limited by the size of RAM, which today is limited to gigabytes in size.

Which Technology Should You Choose?

To determine the best database technology for your business analytics project, you'll need to consider the number of data sources, the need for ETL, the complexity of the data, and the scale and scope of the project.

For projects with limited scope that utilize a single data source, a data mart solution is probably your best bet. When your requirements grow to include multiple data sources with terabytes of data, along with your data analytics needs, a data warehouse is the solution to support scale.

Data Warehouse - When?	Data Mart - When?
<ul style="list-style-type: none"> ▶ For a single centralized data store to serve multiple users and multiple business scenarios (single version of the truth) ▶ When data volumes are large, rapidly growing or may unpredictably spike 	<ul style="list-style-type: none"> ▶ When there is only a single data source, which means the data doesn't need to be consolidated (ETL) prior to being delivered for business analytics ▶ When there aren't many different attributes and metrics to cross-reference (the data mart doesn't need to have many fields) ▶ For a one-time project (e.g. one dashboard), with no added requirements, new data sources or other changes expected in the future

3. What Do I Need To Consider When Selecting A Vendor?

When you begin your research for a BI solution, you will quickly take note of the sheer number of vendors out there. You must narrow the overcrowded field of vendors in order to come up with your own BI short list.

Exploring vendors' websites is usually a great place to start your review. Based on what they share on their websites – and what they don't – you can create a comparison map to quickly figure out which vendors are right for you and which ones probably aren't. Armed with your short list, you can begin contacting vendors to receive proposals.

The Customers Page

A BI vendor website without a dedicated page of featured customers is cause for concern. Assuming the vendor does indeed have a 'customers' page, the first thing you should look for is whether the featured customers are big corporations or smaller companies.

This is an important distinction because BI solutions for big corporations often have very different functional requirements than BI for smaller companies have. You should focus on vendors that sell to companies similar to your own.

If you need BI for an SMB, a BI vendor that lists only Fortune 500 corporations on its 'customers' page probably won't fit your needs as its solutions are probably too complex and/or expensive. Similarly, if you need BI for a large corporation and the business intelligence vendor lists only SMB customers, the solution may not deliver the functionality, performance or scalability you need.

The second thing you should look for is whether you recognize any of the customers listed on this page. Having recognizable names as clients says a lot about the credibility of the BI vendor. Well-known companies with recognizable brand names don't trust their business operations to just anyone. A BI vendor with recognizable names on its 'customers' page is less likely to disappoint than a company listing unknown names.

Online Case Studies

Reviewing the case studies on a vendor's website is a good way to determine whether the solutions the vendor provides were implemented in business scenarios that resemble your own. Credible vendors should include at least a few case studies on their website.

It's important to note whether the solution has been implemented in companies within the same industry as your own. Different industries will share many common needs, but there will also be special needs and requirements unique to each one.

Another important thing to look for is whether the solution is being used in the same way you'd like to be using your BI solution.

Several types of BI solutions are available – some focus on reporting, while others focus on analytics. Some are maintained primarily by IT, and some are designed for more hands-on use by business professionals. Case studies can help you determine whether the solution can meet your goals and whether it can work for the intended users.

The Partners Page

The 'partners' page is a great place to gather information about the vendor. You may think that the 'partners' page is used mainly to find representatives in your own region, but there is a lot more you can tell about a vendor from this page.

One thing you should look at is whether the vendor has a long list of service integrators listed as partners. This is important because it says a lot about how difficult the solution may be to implement. While software vendors prefer to sell software licenses, service integrators typically make their revenue from projects they execute, charging by the hour or by the day.

When a BI vendor has a long list of service integrators as partners, it usually means the vendor's solution is a challenge to implement and/or requires extensive, ongoing maintenance and consulting. Otherwise, the service integrators wouldn't see much business value from working with that vendor.

The second thing you should look for on the 'partners' page is whether the vendor has multiple software/technology vendors listed as partners. Typical BI applications require several tools and technologies to be fully implemented, and when a vendor lists software/technology partners, it usually means the vendor provides only a portion of the BI stack itself.

Download a Self-Service Free Trial

Even though free trial downloads are commonly found in the software industry, they are not commonly found on BI vendor websites. There are two major reasons for this: The first is that 99% of BI vendors are not software providers but rather service providers. As opposed to selling products, they partner with existing software/technology vendors to use those vendors' products as building blocks for implementing custom solutions.

The second reason vendors don't usually have a free trial available for download is their lack of confidence in your ability to quickly and easily use their software. Generally these solutions are too complex or require professional services to get started.

When a vendor does have a free trial version available for download, it typically means that the vendor is willing to sell directly to you (thus removing the markups taken by third-party service providers and resellers) and that this vendor believes you will be able to use its software yourself.

The presence of a free trial also stands to reason that the solution has shorter implementation times without extensive third-party service provider fees.

Professional Services

Vendors in the business intelligence space have two ways of making money: selling software and selling services such as consulting, implementation and maintenance.

It is very difficult for a vendor to successfully focus on both, because they require completely different core competencies, staff and operations. Building commercial software in the BI space requires significantly deeper pockets than simply providing services that utilize existing or third-party software. This is why the BI space has very few vendors that provide end-to-end 'productized' software and significantly more service providers that will custom-build a system.

It's important to understand that the total cost of owning a BI solution is rarely limited to the cost of software licensing but, rather, includes the cost of all the work that needs to be done in order to customize this software for your own needs.

If a BI vendor presents itself as a software vendor but has a wide range of professional services for sale, this should make you wonder what it is about the vendor's software that requires so many services.

As previously mentioned, true software product vendors prefer to focus on selling licenses, not on selling services. It is very difficult to do both equally well, and from a vendor's business model perspective, it's often wrong to even try. BI software vendors try to make their software as easy to implement and as self-service as possible. Otherwise, they would have to spend more than they want on training, support and customization services.

The Pricing Page

A 'pricing' page is not something you often find on BI vendor websites. Those that do have a 'pricing' page usually will not state actual figures but rather a prompt to contact the company for a quote.

The fact that a company doesn't "name its price" on its website does not necessarily mean there's a red flag. On the contrary—most BI solutions aren't of the *one size fits all* variety, as the required service and consequent costs to the service providers vary greatly according to the amount of data that needs to be processed as well as its complexity.

Additionally, some BI vendors do not publish pricing information because most of the actual cost of system deployment will go toward the implementation and customization projects. These projects are very hard to price before doing extensive on-site analysis and estimating project scope. Therefore, standardized pricing would be meaningless.

Some vendors that actually sell software products will refrain from posting pricing information to hide the markups its service integrator partners get on the deal.

In many cases, the service integrator is the one doing the selling to the end customer, and these partners tend to enjoy significant margins—another reason to steer clear of vendors who don't develop their software in-house.

Online Community

Having an online community via blogs and forums is a very big benefit for customers in the BI space as it allows customers to leverage the experience of other users to accomplish tasks and solve problems. Regardless of which vendor you choose, you will always have questions, and it's great to be able to get answers without having to rely on the vendor's official support channels alone.

Online communities also give you a good indication of how popular the solution is and what types of things people are doing with it. Generally speaking, your experience with a BI solution will be much more fluid if you have access to the collective knowledge of an active online community.

When a BI vendor does not have an online community, you should immediately be concerned about how much support will cost you and how quickly you will be able to get responses to your inquiries.

4. How Do I Launch A Successful Proof-Of-Concept?

With your understanding of the software stack and different technology platforms, you should develop a very short list of vendors (no more than one or two) to now to launch a Proof-of-Concept (POC). These tips will put you on the right track for conducting a successful business intelligence POC.

A POC That Doesn't Use Your Data, Doesn't Prove Much

From the vendor's perspective, the most challenging part of any POC is getting familiar with the buyer's data sets and dealing with the unique challenges of that data.

Therefore, most vendors will try to limit the scope of the POC by reducing the scope of data sources involved, trimming the data down, using sample data sets instead of your own data, and so forth.

Since dealing with your unique data challenges make up 80% of the ongoing technical challenges you will be facing throughout your deployment, taking shortcuts here really means the vendor's POC has proven nothing.

Do Not Get Distracted By Pretty Visuals

The visualization of data is important, of course, but the biggest mistake buyers make is to judge the vendor based on the pretty dashboard samples it shows you on its website or during a demo. With visualization software components a dime a dozen, a vendor can easily 'fake' these pretty graphics.

The real challenge is customizing these dashboards to your own needs and having them show your own data. This part can take most vendors months and thus can end up eating a large percentage of your budget. If a vendor can't get your own data to display in dynamic and creative ways within just a few days, you should probably find a better one.

Address Both Future & Present Requirements

Most vendors insist that you define your requirements very clearly to ensure the success of a project of any scope, POC or otherwise. The catch is that it doesn't matter who you are or how much experience you have – it is almost impossible to know in advance what your future requirements will be.

BI requirements tend to be highly dynamic because businesses change all the time and business users are continually refining and adjusting their requirements. Today's reporting needs will look very different in a year from now, and today's analysis will likely be relevant for only a short period of time before becoming obsolete.

The best advice in cases like these is to focus the POC on the ability to meet ad hoc requirements, not just a set of fixed reports that might end up becoming irrelevant by the time the project ends.

Consult Your Own IT Professionals

In many organizations business analytics solutions are already set up that are highly reliant on IT. Business professionals, frustrated by not being able to extract relevant data quickly and independently, will look for BI solutions that cut IT completely out of the loop, giving them the ability to manage his or her own analytics processes independently.

Nonetheless, it is still highly recommended that you consult your organization's IT professionals regarding topics with which they are more familiar: scalability, integration cycles and so forth. Remember, the vendor has no real interest in sharing issues that can trip up the success of any well-intentioned project.

You Should Not Be Required to Make Significant Financial Investment

Some vendors will promise you the world, but will demand significant upfront investment to prepare the projects, hardware and software you need for a POC. Don't agree to this. Demand at least one solid report or dashboard running over your own data before you agree to any financial commitment.

If a vendor is not willing to work with you before you put money down, it's probably because that vendor would have to spend weeks on development before it can reach that point. That typically means the vendor is either trying to sell you archaic technology or simply trying to pull the wool over your eyes.

Next Steps:

- ▶ [Try](#) Sisense for FREE
- ▶ [Join](#) a Sisense analytics expert for a Demo
- ▶ [Contact us](#)