

2010 WHITE PAPER

Full-Featured Business Intelligence for the Cloud

A Review of the Birst Business Intelligence Suite for Dashboards, Reporting and Analysis



Executive Summary	3
Introduction	4
BI Suite Architecture	5
Design Philosophy	5
Tiered Architecture.....	5
Unified Metadata Model	6
Performance and Scalability.....	6
Security.....	7
Virtualization.....	8
BI Platform	9
Data Connectivity Tier	9
Birst Connect for RDBMS.....	9
XMLA Connector for OLAP Cubes	9
Salesforce.com Connector	10
Google Analytics Connector	10
Live Access for Access to Data in Place	10
Data Services Tier	11
Automated Data Warehousing	11
ETL Services.....	12
Analysis Services Tier	12
ROLAP Engine	12
Advanced Analytics Engine.....	13
BI Capabilities	14
Interactive Dashboarding.....	14
Ad-Hoc Query and Pixel-Perfect Reporting.....	16
OLAP Analysis	18
Conclusion	19

Birst, Inc.

153 Kearny St., 3rd floor, San Francisco, CA 94108

www.birst.com | Email: info@birst.com | Toll Free Phone: (866) 940-1496

Executive Summary

Birst provides the first complete Business Intelligence (BI) suite that is purpose-built for the cloud, combining the economic benefits of the cloud with the freedom to choose where you host your data: on site or in the cloud. Traditional business intelligence solutions require expensive software purchases, long implementation cycles, and costly maintenance. Unfortunately, they also offer disappointing results, since the high cost and inflexibility severely limits the system's ability to serve ever-evolving business needs effectively. As a consequence, complete BI solutions are limited to the few large companies that can afford it.

The Birst Business Intelligence suite was designed and developed from the ground up to satisfy the requirements of all types of organizations, small and large. Thanks to patent-pending technology for automating traditionally time- and resource-intensive tasks in the BI application lifecycle, Birst provides an unprecedentedly low total cost of ownership (TCO) and fast time to value. Available from a single product interface, the Birst BI suite offers complete Business Intelligence functionality for highly interactive dashboards, pixel-perfect banded report writing and powerful OLAP-style ad hoc analysis.

The purpose of this paper, intended for an IT, developer and administrative audience, is to provide a technical overview of the Birst BI suite including a description of the design philosophy, platform architecture and components and the BI suite capabilities

Two thirds and more of the total cost of a BI solution over three years is staffing-related

Introduction

The demand for business intelligence is strong and growing. For the past several years, BI has been repeatedly named as the top priority by Chief Information Officers (CIO) in Gartner's annual CIO survey. Putting the right information in the right hands at the right time is a clear imperative for companies of all sizes across all industries. Yet the promise of enabling fact-based decision making remains elusive for many organizations. Too often BI implementations are characterized by high project failure rates and are falling short in addressing the growing need for pervasive and cost-effective BI.

A typical BI implementation using traditional on-premise tools requires the integration of software and hardware from multiple vendors. The process of converting raw data into useful and actionable insights suitable for business decision making involves significant investment throughout the BI process. The acquisition, implementation and maintenance costs quickly add up to levels where BI becomes feasible only for those few companies who have the requisite budget and resources. According to IDC, between 60% and 86% of the total cost of a BI solution over three years is staffing-related¹.

Birst provides business intelligence as a service. Unlike traditional and Open Source BI platforms, Birst was designed to radically reduce overall costs by providing a fully integrated BI suite that automates much of the manual configuration required by other tools. Birst was developed as an integrated solution by one development team with one goal: enabling faster, easier BI deployments. With a Software-as-a-Service (SaaS) solution like Birst, hardware and maintenance costs are amortized across many clients and thereby greatly reduced. Through the use of automation, Birst also significantly lowers the staffing costs required for a BI project for both the initial build and ongoing development and maintenance.

¹ IDC. "Demonstrating Business Value: Selling to Your C-Level Executives." Three Year Server TCO. Based on more than 300 interviews conducted across numerous platforms, presented in composite form. April 1, 2007

Birst automates many of the time- and resource-intensive tasks in a typical BI project

BI Suite Architecture

Design Philosophy

The Birst Business Intelligence suite was designed and developed from the ground up to meet the varying needs and requirements of organizations of all types and sizes. All its components were developed in house and are tightly integrated, sharing a common metadata layer.

Two primary design principles have guided and continue to guide the development of the Birst BI suite:

- 1. Automate all the steps involved in the development, deployment and ongoing maintenance of a BI solution that can be effectively and efficiently automated.** Through patent-pending technology, Birst automates many of the tasks and operations that are repetitive in nature and do not require user input or control. As a result, Birst allows developers, administrators and end users to focus on value-added tasks.
- 2. Provide an intuitive self-service interface without sacrificing functionality or flexibility.** Birst provides greater BI functionality to a broader range of users. It shields average users from feature complexity, while providing power users with easy access to the full breadth and depth of functionality.

Because of these drivers, the Birst is able to provide highly intuitive, full-suite BI while delivering an exceptionally low TCO.

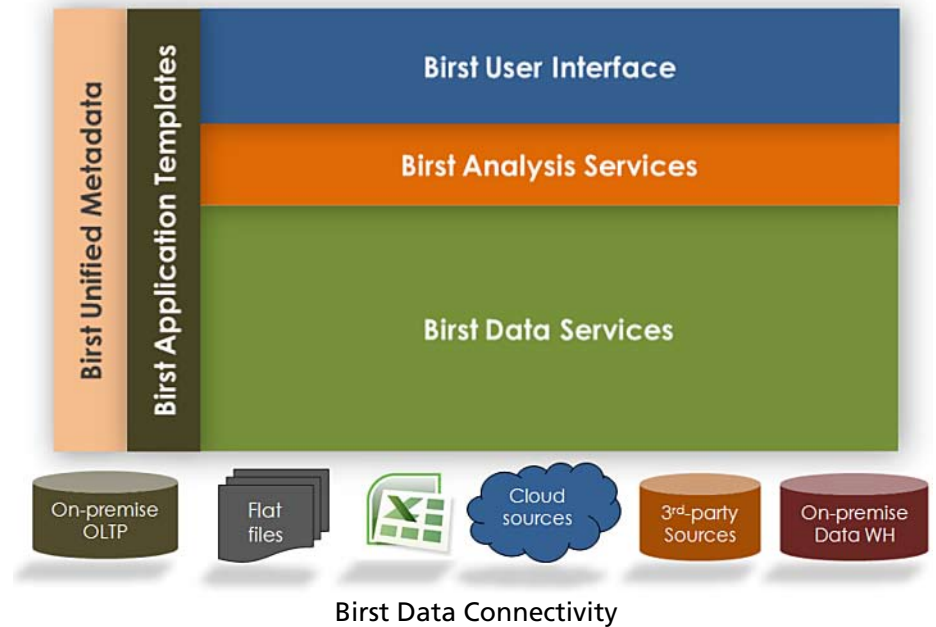
Tiered Architecture

The Birst suite is built on a modern, service-oriented architecture using the latest web standards. The architecture comprises four primary tiers:

- Data Connectivity (includes connectors to on-premise and cloud-based sources)
- Data Services (includes ETL and Data Warehousing)
- Analysis Services (includes ROLAP engine and Advanced Analytics)
- User Interface (includes Dashboards, Reporting and Ad Hoc Analysis)

Each of these tiers will be described in more detail in subsequent sections. The Birst unified metadata layer, described in the next section, underlies these tiers. Application Templates, described in a later section, leverage various components in each tier to enable out-of-the-box, end-user focused analytic applications.

Birst provides high performance and scalability through its shared-nothing architecture and dynamic caching layer



Unified Metadata Model

The Birst Business Intelligence suite architecture is based on a unified metadata model (UMM.) A single, shared layer of metadata ensures consistency across the entire BI suite and eliminates the need to integrate and maintain multiple metadata versions when using a collection of different BI products. Definitions of business-relevant metrics and attributes can be managed in a single place, thereby facilitating the administration of this layer.

Performance and Scalability

Birst is a fully multitenant solution from both a data processing and data storage perspective. Birst's web and application server tier is multitenant meaning that users are spread across an infinitely scalable pool of computing resources without requiring the provision of new infrastructure as usage scales. Leveraging its shared-nothing architecture, Birst is the only BI solution to provide this level of scalability.

Birst has a unique caching layer to help boost performance. In addition to exact and fuzzy cache matching, Birst generates dynamic cubes. These dynamic cubes are indexed structures that provide far better reuse and generate lower database load than traditional caching approaches. The cubes are dynamically partitioned amongst servers to minimize I/O contention and to allow better memory caching, ultimately resulting in a far higher performance solution.

Ensuring that customer data is transferred and stored securely is a top priority for Birst

Security

Providing the right level of data security is critical to a successful BI deployment. As such Birst provides three ways to secure access to data. At the most granular level, Birst provides data-level security via data security filters which gives the administrator the ability to restrict access to specific rows in the data set based on a highly flexible and robust model that allows control over both dimensional and measure data. This data level security can also be inherited directly from the source data during the load process.

Birst also provides the administrator with the ability to control access to the BI assets such as reports and dashboards. This is set up using Access Control Lists (ACLs) which the administrator uses to identify which users get to see which reports and dashboards. Users can also create their own private folders that only they and the administrator has access to.

The third way to define security is by using Custom Subject Areas which gives administrators the ability to craft customized views of dimension columns and measures. This feature is particularly useful to make the application more accessible to business users as well as restricting access to certain data elements.

Finally, as Birst is a SaaS solution it enforces the highest levels of data center security. As such, the Birst service is deployed in accordance with SAS No. 70 that employs stringent controls and processes to ensure complete data security. Birst uses the leading hosting provider, Rackspace, for its data center. Birst's SAS 70 Type II environment includes the following controls:

Physical Security

- Data center access is limited to data center technicians
- Biometric scanning for controlled data center access
- Security camera monitoring at all data center locations
- 24/7 onsite staff provides additional protection against unauthorized entry
- Physical security audited by an independent firm

System Security

- System installation using hardened, up-to-date OS with security patches installed
- Redundant firewalls to help block unauthorized system access
- Data protection with managed backup solutions
- Network-based intrusion detection systems (IDS)
- Continuous risk assessment and security consultation by security professionals

Operational Security

- All employees trained on documented security and privacy procedures
- Access restricted to authorized personnel only, according to documented processes
- Systems access logged and tracked for auditing purposes
- Documented change-management procedures
- Independently audited disaster recovery and business continuity plans for data center

Virtualization

Birst is set up so that administrators can quickly and easily create a new deployment instance with a few clicks of the mouse. For example, a development environment can be moved to production by simply copying the Birst development “Space” and putting the copied Space into production. The ability to easily switch between multiple environments without having to provision new hardware and storage or to install new software also helps in other ways. For example, it helps avoid unnecessarily burdening the production environment with large data loads. The loading can simply be done on a copy of the production environment which can then be switched over to production once complete.

Versatile data connectivity and extract options exist for on-premise databases, flat and structured files as well as popular cloud applications

BI Platform

In this section we detail the features of the platform-level tiers of the Birst BI suite. These are the features found in the Data Connectivity, Data Services and Analysis Services tiers. In the next section, BI Capabilities, we describe the features in the User Interface tier.

Birst's Data Connectivity tier provides options to connect to both on-premise and cloud-based data sources. The Data Services tier delivers ETL capabilities as well as a patent-pending data warehouse generation capability which automates the creation of an optimized star schema from a defined logical model. Its Analysis Services tier includes a ROLAP engine that achieves the highest levels of query performance thanks to multiple layers of caching. It also integrates a predictive modeling and optimization engine that offers advanced analytical capabilities.

Data Connectivity Tier

Birst connectors provide data connectivity and extract options for on-premise databases, flat and structured files as well as popular cloud applications. Birst allows you to extract and upload raw data into its analytically-ready data store. In addition, Birst can also connect in real-time, via Live Access, to an existing on-premise database and provide complete BI capabilities in the cloud without having to extract or upload any data.

Birst Connect for RDBMS

Birst Connect allows you to connect to and extract data from all major relational database management systems (RDBMS) – Oracle, SQL Server, DB2 and MySQL – as well as import entire Microsoft Access database files, Excel spreadsheets, delimited flat files and ODBC data sources. Birst Connect is a Java application that is launched directly from a Birst web client.

Once connected to a relational database management system (RDBMS) or ODBC data source, Birst Connect extracts entire database tables and views, and enables custom SQL statements for custom data extracts. IT can schedule data extraction and processing tasks using either a built-in scheduler or a preferred task scheduler running on local servers.

XMLA Connector for OLAP Cubes

Birst can connect to Hyperion Essbase and Microsoft Analysis Services cubes via XML for Analysis (XMLA.) Birst automatically generates the metadata needed to enable a fully querying of XMLA cubes. Using Birst's data federation capabilities, cube-based data can be combined with other data sources for new types of analysis.

Live Access makes it possible to leverage your existing data warehouses and marts without the need to extract data out and store it in the cloud

Salesforce.com Connector

It only takes Birst a few minutes to connect to and extract data from salesforce.com. The Birst connector for salesforce.com supports both standard and custom salesforce.com data objects.

Unlike other BI solutions for salesforce.com, Birst does not limit reporting and analysis to just salesforce.com data. With Birst you can easily combine your salesforce.com data with other data sources, for example booked revenue from an ERP system or lead data from a marketing automation solution, for more powerful and holistic insights.

Google Analytics Connector

You can experience the full breadth and depth of Birst's reporting, ad-hoc analysis and dashboarding capabilities on your Google Analytics data. The Birst connector for Google Analytics provides a live connection to your Google Analytics account, providing real-time insights into your AdWords campaigns and more. Birst goes beyond the reporting capabilities provided in Google Analytics by providing features like pixel-perfect reporting, automated report scheduling, interactive dashboarding and the option to combine Google Analytics data with other data sources.

Live Access for Access to Data in Place

Birst Live Access allows is a major breakthrough for the SaaS world, enabling IT to leverage the economic benefits of the cloud without having to move data to the cloud. Thanks to this patent-pending technology, IT can connect directly to on-premise data warehouses and use Birst to build reports and interactive dashboards in the cloud without having to move any data offsite. Thanks to sophisticated federation, on-premise data from the Live Access connector can be combined with cloud-based data from other sources for reporting and analysis.

Live Access addresses the compliance and data governance concerns that some IT departments will have regarding SaaS-based solutions. Thanks to Live Access, Birst serves as a scalable ROLAP front-end that provides dashboarding, pixel-perfect banded reporting, and OLAP-style analysis capabilities on the data right behind your firewall. Live Access can connect in real-time to MS SQL Server, Oracle, DB2, MySQL and ODBC data sources. Queries and result sets are transferred securely via SSL.

Birst automatically generates a fully optimized physical database schema from a logical data model

Data Services Tier

Birst Data Services includes a full Extract, Transform and Load (ETL) capability that is tightly integrated with the rest of the suite's BI capabilities as well as a patent-pending automated data warehouse generation capability.

Automated Data Warehousing

Birst makes it possible to extract, transform and load data into a versatile data warehouse design in a fraction of the time required when using conventional BI tools. Thanks to Birst's patent-pending automation features, IT is finally in a position to quickly serve the analysis and reporting needs of business users without having to give up control over the data.

Using conventional BI tool sets, designing and creating a data warehouse is a resource-intensive and time-consuming process. The process of creating and maintaining the data warehouse database schema and mapping the source data to it has traditionally been done manually. Birst on the other hand automatically generates and maintains all the fact tables, dimensions, joins as well as all the required routines for loading data into the analytical data store. The advantage of automation is that it enables the rapid iterative development of a data warehouse by enabling IT to focus on the logical model underlying business processes and not the physical instantiation of it.

Once the logical dimensional model has been defined, Birst automatically compiles the model into a fully optimized star schema. Logical measures are automatically analyzed for calculation grain and logical dimensions are analyzed for levels requiring persistence. Fact tables are automatically generated, while dimensions and mini-dimensions are created as needed. Birst generates and manages all key relationships, including surrogate keys, when necessary.

Full and incremental data loading strategies are available automatically. No additional scripting is required for an incremental load and changes are detected automatically. Birst also provides for automatic management of historical data including the automated support for dimensional attribute history and snapshots. No additional coding is required to leverage these capabilities. Load optimization allows for calculations to be grouped in the same loading pass. Such load groups enable segmentation and parallelization of the loading process. Each load group can have its own periodicity (e.g. day, week, month) and own specific content (e.g. organization). Time-shifting variants (e.g. month ago, quarter ago, etc.) and rolling-time variants (trailing 3, 6, 12 months, year-to-date, etc.) are automatically supported and do not need to be separately set up.

The Analysis Services tier of Birst includes a ROLAP engine and advanced analytical capabilities

ETL Services

For cases where source data cannot be directly mapped onto a desired logical dimensional model – either because the source data structure is too different from the target data structure or complex data manipulations are needed to ensure the data maps correctly – Birst provides an ETL capability. This capability is fully integrated with the rest of the BI suite.

ETL Services was designed explicitly to provide a high level of flexibility. It is based on a fully scalable, in-memory middle tier that does the transformation processing. It stages intermediate tables and performs transformations on data prior to loading. Users can write procedural scripts across records as well as one-by-one and can create complex manipulations like sequence analysis or hierarchy tree navigation.

Analysis Services Tier

Birst's Analysis Services tier includes a modern ROLAP engine that delivers high query performance thanks to smart caching. It also integrates a predictive modeling and optimization engine that offers advanced analytical capabilities.

ROLAP Engine

Birst provides full ad-hoc analysis capabilities without the need for any OLAP cubes thereby offloading IT from the resource-intensive and time-consuming task of constantly having to maintain and optimize a growing cube farm. With Birst, IT simply defines a logical model and Birst takes care of generating an optimized physical instantiation which gives analysts true ad-hoc access to the entire data warehouse. Analysts can easily issue their own ad hoc queries and use drill, filter, and pivot to analyze the result sets, all from a single user interface.

Unlike other OLAP products, Birst does not restrict dimensional access to the data. Birst constructs a dynamic logical cube of all data that it is mapped to. It provides the full richness, scope and depth of information that can be possibly analyzed. As long as dimensional relationships exist between the various data elements, users can analyze the data any way they want.

Birst has a unique caching layer that provides a performance improvement over traditional OLAP solutions. In addition to exact and fuzzy cache matching, Birst generates dynamic cubes to help with performance. These dynamic cubes are indexed structures that provide far better reuse and generate lower database load than traditional caching approaches. The dynamic cubes are dynamically partitioned amongst servers to minimize I/O contention and to allow better memory caching, ultimately resulting in a far more scalable solution.

Birst's ROLAP engine also works on operational tables as well as on star schemas. Birst can map onto "opaque views" which are essentially inline views with multi-pass SQL operations. When combining data at different levels of aggregation, Birst will push down multiple sub-queries and use a multi-pass approach to pull the individual results together into a single result set.

Advanced Analytics Engine

The Birst platform contains advanced predictive analytics and optimization capabilities. In contrast to conventional data mining environments, data does not have to be moved out of the analytic data store and reformatted. Data sets for model training and scoring are generated directly from the analytical data store. Additionally, data set inheritance enables the efficient reuse of a data set to meet multiple predictive modeling objectives.

Birst's advanced analytics capabilities leverage the ROLAP engine for data preparation. The modeling engine makes full use of aggregates and derived measures. Sophisticated new measures are defined and calculated on the fly as inputs into the modeling process. Share, time-series and dimensional breakout metrics are used to enrich the information contained in the data warehouse. The use of OLAP-style measures for modeling enables the addition of complex and highly predictive behavioral calculations.

For each modeling task, Birst automatically evaluates a comprehensive set of data mining techniques. Supported techniques include linear and logistic regression, decision trees, feed-forward neural networks, support vector machines and rules/regression trees. Modeling scores are written directly back to data warehouse tables, ready to be used in ad hoc queries and dashboards or to be fed into additional processing, for example list generation. Both rules-based and model-based recommendations can be combined into complex decisions.

In addition to predictive modeling, the Birst BI platform offers unique capabilities for performance optimization. For a given business entity, the optimization engine determines an optimal set of improvement actions based upon analyzing the impact and probability of a universe of candidate actions. The engine automatically traverses the full dimensional model to identify patterns and relationships that are critical to understanding and maximizing business results, providing analytics that prescribe specific actions to individual business users.

Birst provides a high level of interactivity and self-service capabilities in the dashboard user interface

BI Capabilities

The Birst User Interface tier contains all the suite's BI capabilities including interactive dashboarding, reporting and ad-hoc query and analysis. All of these capabilities are available from a single user interface and are tightly coupled through the unified metadata layer.

Interactive Dashboarding

Birst's Interactive Dashboards provide an unprecedented level of BI self-service from an easy-to-use interface designed for business users of all types. Birst goes beyond the old dashboard paradigm of providing static "at-a-glance" information on business performance for managers and executives. With Birst, users can directly interact with the dashboard data without any formal training or specialized BI expertise.

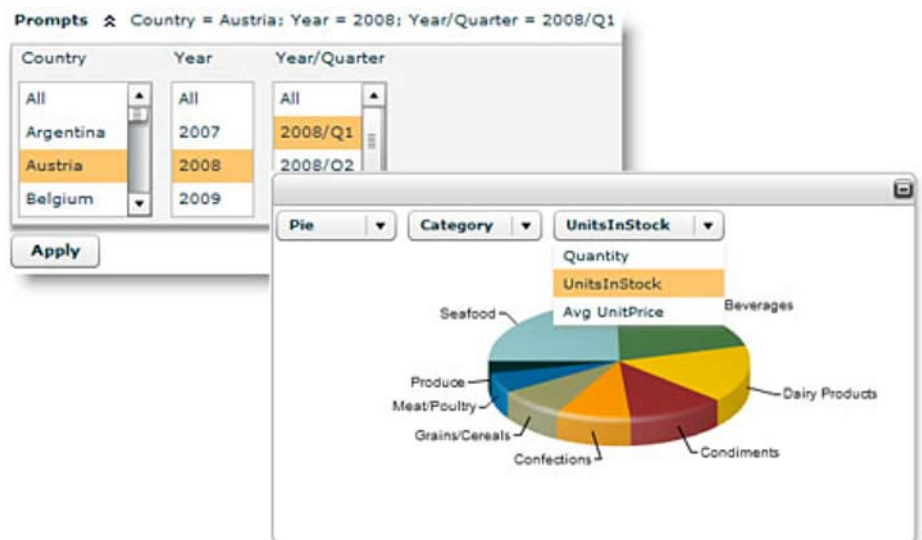


For example, if the users want to look at a different combination of metrics and attributes they can simply select from dropdowns in the dashboards. If they need to visualize a given report as a pivot table they just drag and drop the desired columns into place.

Users can filter the results on the page via filter prompts, by clicking on segments in a chart, or by drawing a rectangle around a set of values in a table.

CategoryName	ProductName	Sales
Beverages	Chai	12,788.1
Beverages	Chang	16,355.96
Beverages	Chartreuse verte	12,294.54
Beverages	Côte de Blaye	141,396.73
Beverages	Guaraná Fantástica	4,504.36
Beverages	Ipoh Coffee	23,526.7
Beverages	Lakkalikööri	15,760.44
Beverages	Laughing Lumberjack Lager	2,396.8
Beverages	Outback Lager	10,672.65
Beverages	Rhönbräu Klosterbier	8,177.49

Birst Dashboards offer drill-down and drill-across capabilities that make delving into report details as simple as point and click. Both charts and tables are easily configured to drill to any desired target column or dashboard. Users can also drill out to other applications or web pages. The “drill anywhere” capability allows users to drill using new dimensions not present in the chart or table.



Creating new dashboards ones is just as straight forward using drag-and-drop functionality. Any Birst report can be exposed in a Dashboard. If the user needs to add prompts, they can just drag and drop the appropriate columns from the subject area onto the dashboard canvas.

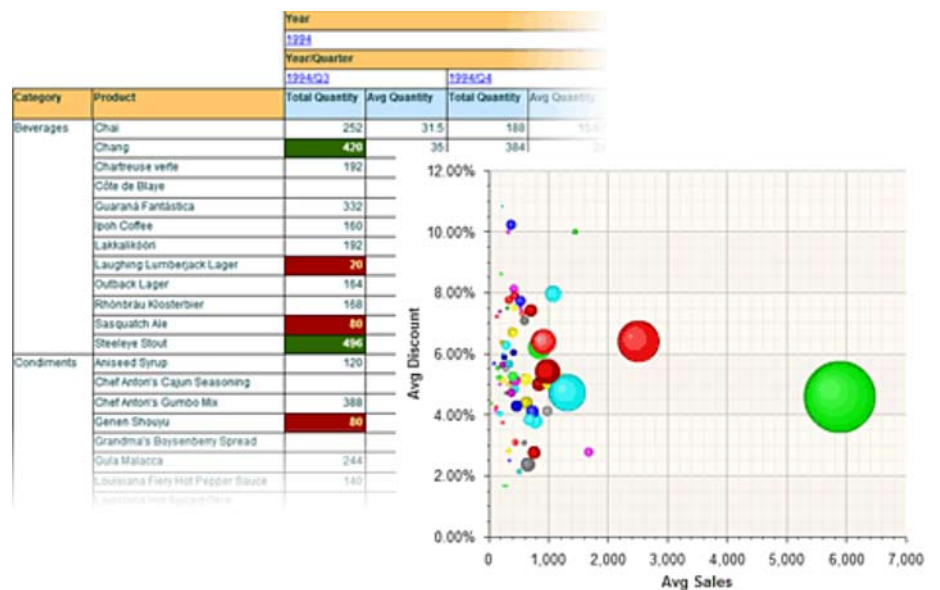
Birst provides a single user interface for ad hoc analysis and pixel-perfect banded report writing

Key features at a glance:

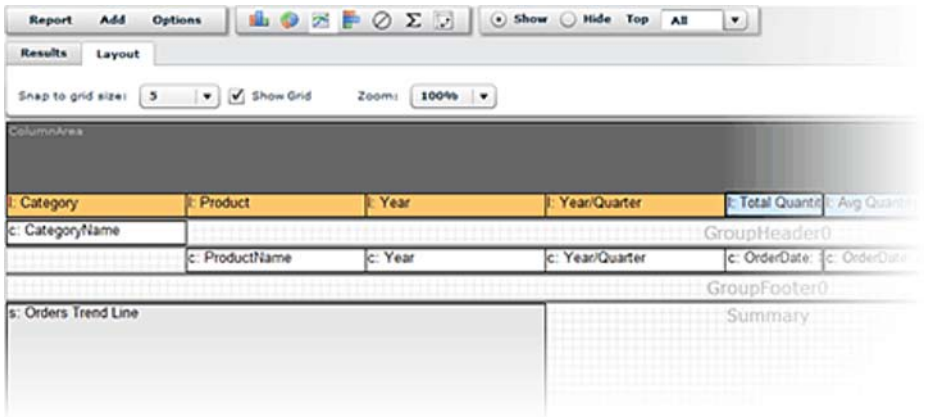
- Ad-hoc analysis in the dashboard UI
- Interactive charts and graphs
- Drill down and across dashboards
- Expand & collapse dashlets
- Cascading dashboard prompts
- Intuitive drag-and-drop functionality to create dashboards

Ad-Hoc Query and Pixel-Perfect Reporting

Unlike other BI tools which require multiple products, Birst provides enterprise reporting, ad-hoc query and analysis in a single easy-to-use interface. Birst enables users to quickly perform ad hoc queries, create pivot tables via drag & drop and visualize data using advanced chart types. Both IT-provisioned and user-created reports can be saved and shared for reuse across the organization and are easily turned into interactive dashboards with only a few clicks.



Every data element in the analytical data store, including modeling results, plus all logical measures can be made accessible for OLAP-style analysis and report generation. Birst also includes a fully-featured report layout editor based on an object-oriented design for advanced report creation. Full extensibility ensures that any conceivable report page can be designed. Reports are compiled into Java byte code for fast and direct execution. No interpretation at runtime is required. Performance is enhanced by server-side processed report caching. Integrated cache-bursting technology ensures scalability and greater mobile functionality.



Key features at a glance:

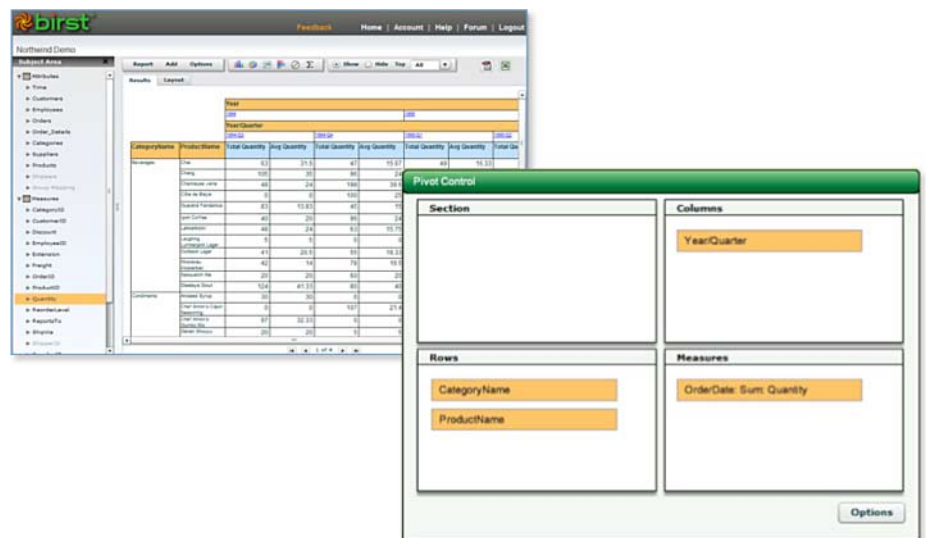
- All capabilities available through a single UI without the need for plug-ins or Java applets
- Feature-rich WYSIWYG editor with one-click toolbars and drag & drop functionality
- Pixel-perfect banded report writing both on screen and on paper
- Powerful ad-hoc analysis using drilling and OLAP-style functions
- Automated scheduling and delivery capabilities for report distribution
- Real-time alerts and conditional formatting for exception reporting
- Report export options to PDF, CSV, Excel and more
- Report caching for optimal performance
- Integrated cache-bursting for scalability and greater mobile functionality
- Wide variety of chart types for data visualization: Bar, Column, Pie, Line, Bar/Line, Area, Stacked, Meter, Bubble, Tree and Heat Maps, Pyramid, Funnel and Geo Maps

OLAP analysis supports positional references, cell-based calculations and differentiation between slicers and filters

OLAP Analysis

Birst provides best of breed capabilities for multidimensional analysis. Quickly slice and dice your data from a number of different perspectives as well as identify complex data relationships through the use of powerful cross-dimensional calculations.

Birst's ROLAP engine provides support for advanced concepts like positional references, cell-based calculations and differentiation between slicers and filters. Using powerful positional calculations, Birst enables users to analyze how a value in a given cell of the cube relates to values elsewhere. Birst's best-in-class calculation capabilities include dimension-specific aggregation, inheritance, business rules, multi-pass calculations and virtual measures.



Birst, Inc.
153 Kearny St., 3rd floor
San Francisco, CA 94108
www.birst.com
Email: info@birst.com
Toll Free Phone: (866) 940-1496

Conclusion

Birst provides the first complete BI suite that is purpose-built for the cloud, combining the economic benefits of Software-as-a-Service (SaaS) with the freedom to choose where you host your data: on site or in the cloud.

- **Complete Business Intelligence**

Birst is an end-to-end BI suite that supports all styles of Business Intelligence in a single product. The Birst BI suite includes all the tiers and components of a modern BI architecture. It rivals traditional on-premise BI vendors in both breadth and depth of functionality.

- **Lowest TCO**

Birst offers much lower total cost of ownership (TCO) compared to on-premise BI solutions. As a pure SaaS solution, Birst drastically reduces both upfront build and ongoing maintenance costs. Thanks to patent-pending technology, Birst automates many time- and resource-intensive steps in the BI solution lifecycle to further reduce staffing costs and provide faster time to value.

- **Host Data in the Cloud or On Site**

Birst combines the advantages of two worlds: the economics of the cloud with the functionality and security of on premise software. Customer data does not have to leave the firewall. Birst is the only cloud BI solution that provides the option to keep data where it currently resides.