A platform for networked business analytics

BUSINESS INTELLIGENCE

Infor® Birst's unique networked business analytics technology enables centralized and decentralized teams to work collaboratively by unifying IT-managed enterprise data with user-owned data. Birst automates the process of preparing data and adds an adaptive user experience for business users that works across any device.

This white paper will explain:

- Birst's primary design principles
- How Infor Birst® provides a complete unified data and analytics platform
- The key elements of Birst's cloud architecture
- An overview of Birst security and reliability.

61% of Birst's surveyed reference customers use Birst as their only analytics and BI standard.[1]

Leveraging Birst with existing enterprise BI platforms

Birst networked business analytics technology also enables customers to leverage and extend their investment in existing legacy business intelligence (BI) solutions. With the ability to directly connect to Oracle Business Intelligence Enterprise Edition (OBIEE) semantic layer, via ODBC, Birst can map the existing logical schema directly into Birst's logical model, enabling Birst to join this Enterprise Data Tier with other data in the analytics fabric. Birst can also map to existing Business Objects Universes via web services and Microsoft Analysis Services Cubes and Hyperion Essbase cubes via MDX and extend those schemas, enabling true self-service for all users in the enterprise.

Contents

Agile, governed analytics in the era of data discovery

The business intelligence and analytics landscape has been experiencing a significant transformation for several years now. With massive volumes of data, more data living outside the enterprise data warehouse, and increasing user demand for speed, autonomy, agility, and smart analytics, organizations are struggling with an increasing divide between end users and centralized IT teams.

End users, driven by a thirst for data-driven daily decisions, have started their own analytic initiatives on decentralized data using desktop data discovery tools. These "shadow" initiatives have increased end-user autonomy and provided instant gratification but have also created analytical silos and dangerous inconsistencies in data analysis.

The centralized teams—gatekeepers of mission-critical data—are burdened with legacy technologies, legacy reporting requirements, and legacy processes, all of which have prevented them from meeting the business's need for speed, agility, and personalized insights.

Without this agility that the business demands, or the consistency and governance IT require, an organization cannot become data-driven. In the 2018 Magic Quadrant for Analytics and Business Intelligence Platforms, Gartner predicted that: "By 2020, organizations that offer users access to a curated catalog of internal and external data will derive twice as much business value from analytics investments as those that do not." [1]

It's clear that mistrust in the data provided by discovery tools results in more arguments over numbers and less time spent making data-driven decisions. It's in this context that Birst provides networked business analytics technology that enables centralized and decentralized teams to collaborate around a "shared version of the truth."

Birst's unique networked analytics technology enables IT leaders to govern, support, and scale multiple integrated environments—while providing end users with autonomy, ease-of-use, and speed to work with non-curated and curated data. This approach allows independent teams to analyze user-generated data blended with governed enterprise data. It also enables the centralized team to better serve their end users by providing true self-service across a single view of all business data for not only the analyst but also for non-data-savvy business users fostering the confidence and trust in data that senior executives demand.

Design principles

Birst was designed from the ground-up based on the idea that trusted, well-governed data is not at odds with speed and ease of use. It leverages new capabilities made available by modern technologies like artificial intelligence and cloud computing—multi-tenancy, virtualization, machine learning, and web-scale architectures—to truly combine the centralized and decentralized models of BI, delivering the best aspects of both: enterprise scale end-user self-service without analytical silos.

At Infor, we believe that solving these real-world analytics problems is something that cannot be fully addressed by the user interface alone. While providing an intuitive experience is critical to analytic success, the most significant challenge in analytics continues to be unifying and refining data for business use—making data "business ready." It could be argued that one of the biggest weaknesses in the dash to data discovery has been that it fails to address the complexity of data in most organizations. Birst embraces that complexity and provides unique solutions to taming it.

Working with data to support analytics should not be the exclusive domain of IT. Business users should also have the tools to prepare their own 'edge' data and combine it with the trusted, centrally managed data from IT to analyze complex business processes. Birst provides both easy-to-use data preparation capabilities for self-service as well as the powerful data integration technologies required by IT.

For everyone to make confident decisions, it is critical to maintain consistency and trust in the data. Four primary design principles create that trust and guide Birst's approach to data management.

1. Data refinement

Birst leverages intelligent unification technologies that both map and model data from multiple sources. Whether it is existing warehouses, data lakes, cloud applications, or transactional systems, all data is captured and unified regardless of size, structure, or speed so that there is one consistent view of the data.

Birst can combine multiple data sources that each have their own definition of "customer" and unify these disparate sources into a single version of "customer" for all users.

Birst leverages its pre-built connectors, live access, a business-ready Data Store, query federation, intelligent data navigation, and a wide range of data mapping and extracting capabilities to accomplish data unification.

Infor believes that all data must be refined before it can be used by business users. The refinement can be as simple as turning 15 operational data tables into a representation of facts and dimensions or as complex as creating a business rule that leverages data from diverse and constantly changing data sources to create a common and reusable business metric.

The historical challenge in data refinement is that it takes too much time and too many resources to refine the data for day-to-day use and managing changes in underlying data structures prevents this process from keeping up with new business demands. In recent years, organizations have been investing in data lakes to remove this overhead.

Although a data lake does remove much of this overhead, they are structures better suited to data scientists and tend to be unsuitable for normal business use. To overcome the speed and data source changing challenges, Infor Birst introduced patented Automated Data Refinement (ADR), a complete extraction ion transformation and loading (ETL) language, and smart data change detection to enable enterprises to create an agile semantic layer—or a "shared version of the truth"—that adapts at the speed of business today.

2. Networked insights

A fundamental capability of Networked BI is multitenancy, which enables the creation of virtual—not physical—BI tenants that relate to each other.

The use of virtual instances is important because, traditionally, delivering trusted and reliable data across the enterprise largely depended on physical replication of BI infrastructure—not just hardware but also data, metadata, user profiles, system configurations, etc.—making it a time- consuming and expensive effort.

To allow business teams to work on their own, while staying networked to a central, governed data set, Birst allows different groups, such as finance, customer support, sales, and marketing to use their virtual copies of the high-performance data tier to gain access to centralized data and blend that with their local data and spreadsheets. This paradigm creates consistency and collaboration between IT and business teams, ensures centralized governance, and empowers data ownership, independence, and self-service data blending at the point of impact.

Networked BI creates unique and exciting possibilities. If we consider the analytical fabric as an organically grown or "crowdsourced" network of insights, it becomes a powerful method for harnessing the collective intelligence of an organization, turning the idea of "enterprise business intelligence" into a reality.

3. Adaptive user experience

Infor Birst's vision is for every individual within an organization to have business data at their fingerprints to improve even the smallest decisions they make. Business people must have tools that provide flexibility and freedom to answer any question and blend their own user-generated data with enterprise data. To do this, Birst delivers a user experience that uniquely meets everyone's needs by supporting all the different styles of analytics, including pixel-perfect reports, highly interactive and responsive dashboards, intuitive visual discovery, native and offline mobile, embedded predictive tools, and self-service data preparation.
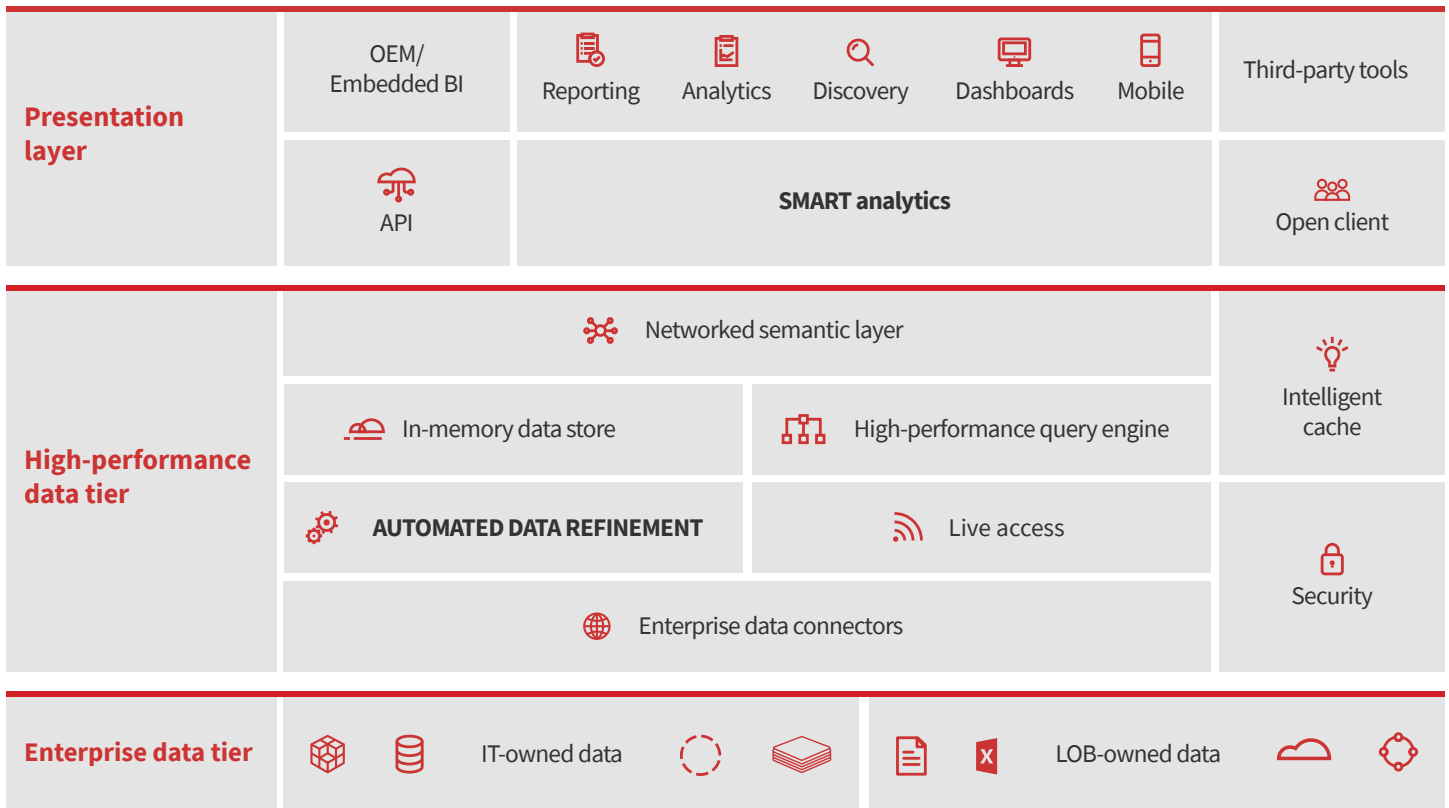
Birst's user experience blurs the traditional lines between dashboards, reporting, and discovery to create an adaptive interface where any user can automatically surface new, personalized insights, with AI-powered BI and easily interact with business data to make better decisions. Equally important, these decisions can be made with full confidence because all analytics styles pull data from Birst's agile semantic layer, which delivers the necessary governance to ensure a single version of the truth.

4. Smart analytics

The key design principle underlying Birst Smart Analytics is to make it easier for the average business user to make data driven decisions. This is realized by applying Birst's patented machine learning algorithms to dashboard KPIs to automatically diagnose or uncover drivers affecting business performance. This leads to machine automation of tasks to accelerate time to insight, discovery of relationships that the human may not have even considered or doesn't have the skills to perform and to move analytics from the descriptive to diagnostic and predictive—Smart Analytics helps business people go from "what happened?" to "why did it happen?" or to "what will happen?"

We build on Birst's existing capabilities to offer AI-driven insights that are designed for a business user and can be personalized yet still governed through Birst's agile semantic layer. The combination of smart analytics and networked BI means that smart analytics can be applied across extended enterprise datasets, not just personal data sources.

Modern enterprise-class architecture

Presentation layer	OEM/ Embedded BI	Reporting	Analytics	Discovery	Dashboards	Mobile	Third-party tools
	API	SMART analytics					Open client

High-performance data tier	Networked semantic layer		Intelligent cache
	In-memory data store	High-performance query engine	
	AUTOMATED DATA REFINEMENT	Live access	Security
	Enterprise data connectors		

Enterprise data tier	IT-owned data	LOB-owned data

These design principles are supported by an enterprise analytics architecture. Birst's native cloud architecture provides a low total cost of ownership (TCO), rapid deployment, automated zero-cost, zero-pain upgrades, programmatic life cycle management, and standards-based, completely open integration. These capabilities enable IT teams to build production-scale analytical applications that deliver business value daily and greatly reduced administrative costs. Furthermore, the benefits of the cloud extend to the business, where a single business analyst can administer an entire Birst deployment.

A networked approach to business analytics

To deliver on the design principles described above, Birst provides a complete Networked Business Analytics platform from data acquisition, to transformation, modeling and analysis—all within a fully integrated, native cloud architecture.

Enterprise data tier

Data connectivity

Birst provides data extraction and connectivity options for a wide variety of databases, data lakes, flat and structured files, in-memory and analytic databases, as well as popular cloud and on-premises applications. Birst supports the extraction of entire database tables or views, and the extraction of subsets of data using custom SQL queries. Birst Connect, a Java™ application sitting on-premises, can be used for both bulk data extraction and for connecting in real-time (Live Access).

Birst extraction tasks can be scheduled either using Birst's cloud-based orchestration as a service feature or an external OS scheduler. In addition to extraction of data from all relational and analytic databases (SQL Server, Oracle®, HP Vertica, Teradata®, Amazon® Redshift, Snowflake, etc.) and applications, Birst supports uploading delimited flat files, JSON, XML, Microsoft® Excel®, and Access database files. Structured data is extracted and uploaded in a tabular format of columns and rows. These extract and live query capabilities also extend to modern (unstructured) and big data sources such as Hadoop Hive, Cloudera Impala, Cassandra™, Google™ Big Query, Infor Data Lake, and more. In all cases, data is transferred securely to Birst using secure authentication, encryption and compression techniques.

For even faster deployments and zero connector maintenance, Birst offers an extensive catalog of pre-built connectors to the most popular cloud and on-premises business applications. These connectors, leveraging web-service APIs and JDBC, have been designed to extract standard and custom objects, or even specific columns, from the respective applications and are maintained over time to ensure connectors stay up to date with application changes. Birst also supports JSON files and REST and SOAP for connectivity to the latest web services.

Birst also offers Live Access (real-time query) capabilities to query on-premises data sources directly, like XMLA cubes, existing enterprise data warehouses, data marts, applications, and data lakes. On-premises data stores or applications are queried in real-time without the requirement to first extract and load the data into the cloud. This helps to bridge the gap between centralized and decentralized teams, enabling enterprises to leverage their existing investments in data warehouses, data lakes, data marts and XMLA cubes (their centralized data assets), while still leveraging Birst for combining decentralized data and creating the essential analytics fabric. Live Access connects to on-premises data sources directly, in real-time, and transfers query results securely over TLS encrypted channels (HTTPS).

High-performance data tier

High-performance analytics without the hassle

Birst automatically compiles a logical, dimensional model into a modern star schema design and generates a physical, fully optimized, star schema without writing any code. Logical measures automatically turn into calculation grains and logical dimensions automatically turn into levels. Fact tables, dimension tables and joins are automatically generated and maintained as are all the required routines for loading data into the Business-ready Data Store. Full and incremental data loading is available automatically.

No additional scripting is required for an incremental load, and Birst also provides automatic management of historical data including snapshots. This approach flips traditional warehousing on its head by leveraging the logical model to create the physical model and is the reason Birst can deliver not only speed, but also governance, because the physical data reflects the business definition of the data and not vice versa. This patented technology is core to the networked approach and brings together the agility of business with the governance of a logical dimensional model.

Automated Data Refinement (ADR)

Birst's cloud analytics engine delivers patented automated data integration capabilities for most data integration needs and a developer-friendly scripting (ETL) language for more complex needs. An example of this automated refinement is Birst's automated time-series measures, where all measures are automatically available by common time- series dimensions, like trailing twelve months, trailing three months, etc. Since Birst is a single integrated platform, all data integration routines are developed, tested, and enabled in production from a single web browser, without any work required in a different application or desktop-based tool.

In-Memory Data Store

Birst's high-performance In-Memory Data Store seamlessly combines different sources of data. It is designed and optimized for ROLAP-style analytics, providing a Kimball-style star schema with a multidimensional view of all data. In addition, Birst supports Type 1 and 2 slowly changing dimensions, conformed dimensions, and manages snapshots and time-based transformations automatically. Data loading and updates are done through incremental processes with built-in change detection.

Birst high-performance analytics engine

Birst's Analytics Server is comprised of three tightly integrated components: an agile networked semantic layer, a ROLAP Engine, and a data navigator. The section below describes how these work together.

Agile networked semantic layer

Birst's patented ADR combines and organizes data from multiple sources into a Business-ready Data Store and overlays it with a common and reusable semantic layer. This is a single set of business rules and definitions that enable data governance and ensure that every user, regardless of who they are or how they access their information, can trust the veracity of the data they're consuming.

Birst's semantic layer enables users to create their own custom measures and attributes while still delivering a shared version of the truth to the entire organization.
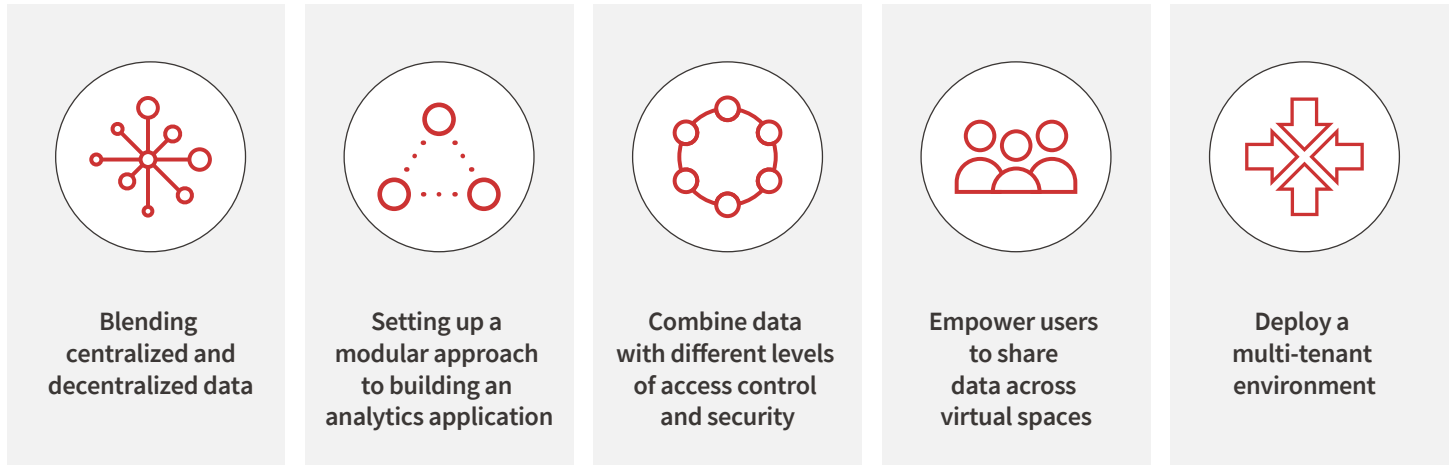
Also, an agile semantic layer significantly eases administrative and development tasks by taking advantage of centralization and reusability, so any changes to underlying data structures are automatically propagated across the environment. The lack of automation with legacy semantic layers was one of the flaws of these early BI platforms. The Birst networked semantic layer defines how the physical Business-ready Data Store is created—which enables Birst's unique networked approach—providing end-user flexibility while maintaining a single unified version of the business, enabling users to work independently, yet build collaboratively.

ROLAP engine—The Birst ROLAP engine provides full, ad hoc analysis capabilities without the need for physical OLAP cubes, thereby removing the resource-intensive and time-consuming task of constantly having to maintain and optimize cube farms. Unlike other OLAP engines, Birst does not restrict dimensional access to the data. Birst constructs a dynamic logical mapping of all data, providing rich and indepth analytics capabilities. The ROLAP engine uses Birst's logical query language (BQL) to enable administrators to query the agile semantic layer.

Data navigator (data modelling)—After attaching to a data store, Birst first maps the metadata within the data store. Mapping ensures that Birst understands the form and structure of the data inside the database, but Birst only extracts data from the source if the user requires it. The Data Navigator decides when to access data in the Businessready Data Store or when to access data on-premises, and even when to access data in both places and combine it within the query itself (Query Federation).

All of this is managed by the Birst data navigator, and the end user does not need to know the specific data elements that create a business metric (like lead conversion rate). Instead, they can focus on analyzing the metrics and decision making, abstracting away the complexity of multi-tiered data models. With the option to have Query Federation of Live Access data sources along with cloud data in Birst's Business-ready Data Store through the agile semantic layer, customers can quickly and easily boost the value of existing on-premises and cloud-based data assets, such as data lakes, by allowing high-performance aggregate data to reside in Birst, but still allowing users to seamlessly drill down to detail data stored on-premises.
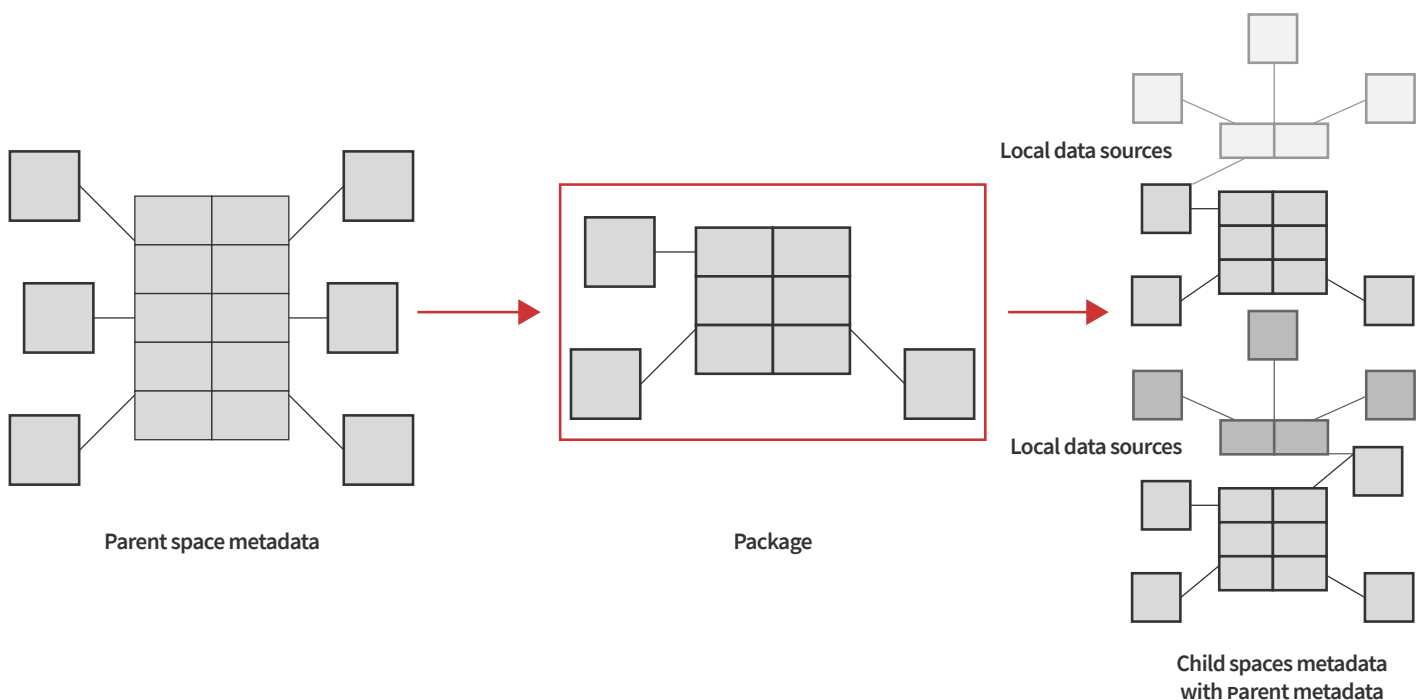
Virtual spaces are useful in many situations, including:

Blending centralized and decentralized data	**Setting up a modular approach to building an analytics application**	**Combine data with different levels of access control and security**	**Empower users to share data across virtual spaces**	**Deploy a multi-tenant environment**

Virtual spaces—Spaces are virtual instances of analytics that can be shared by multiple users and contain uploaded and processed data, the corresponding metadata, as well as all reports and dashboards for an analytics use case.

Users can create an unlimited number of spaces for different data sources or analytical purposes, creating a true multi-tenant solution for every enterprise. Spaces are at the core of Birst's networked business analytics approach. They are the foundation of the shared version of the truth, allowing centralized and decentralized teams to leverage and extend each other's insights.

Centrally provisioned packages

Parent space metadata

Package

Local data sources

Local data sources

Child spaces metadata
with ᴘarent metadata

Centrally provisioned packages

A package is a set of metadata that can be imported from a parent space to a child space, augmenting the existing metadata in the child space in the process. This networking effect allows users to have access to data from multiple sources by blending data residing in one space (parent) with data in another space (child). Users assemble metadata in the parent space into packages that are then shared with child spaces.

A key design principle here is that no data is moved between spaces. Instead, data is virtualized and extended logically This way the child space inherits changes from the parent space automatically and any risk of data inconsistencies are eliminated.

Another key design principal is that nothing in the child space can alter the parent space. For example, if metadata is exported from a space to be used by other spaces, that parent space is not affected by anything that happens in a child space.

Self-service data preparation

Self-service data preparation enables governed data discovery, bridging the gap between IT's charter to govern and maintain centralized data and the need for business to add new data for local analysis flexibly. Freeing IT from having to provide individual data feeds to departments and individuals, self-service data preparation provides business users with personal, analytic sandboxes while ensuring data security and compliance. This allows for independent development of separate subject areas, while logically connecting them to create a shared version of the truth that can blend at enterprise scale and preserve rapid independent development.

Birst Connected Data Prep takes the power of ETL and puts it in the hands of business people with a user-friendly interface using a very simple three-step process:

- **Connect** to different data sources.
- **Prepare** your data using powerful transformations.
- **Relate** your data with the rest of the Birst network of analytics.

By leveraging Birst Spaces, administrators can provide a sandbox environment where business users can confidently add new data and conduct their own analysis. New data, reports and dashboards created by the business user can then be promoted to the enterprise environment, seamlessly incorporating them into the agile semantic layer to maintain data governance.
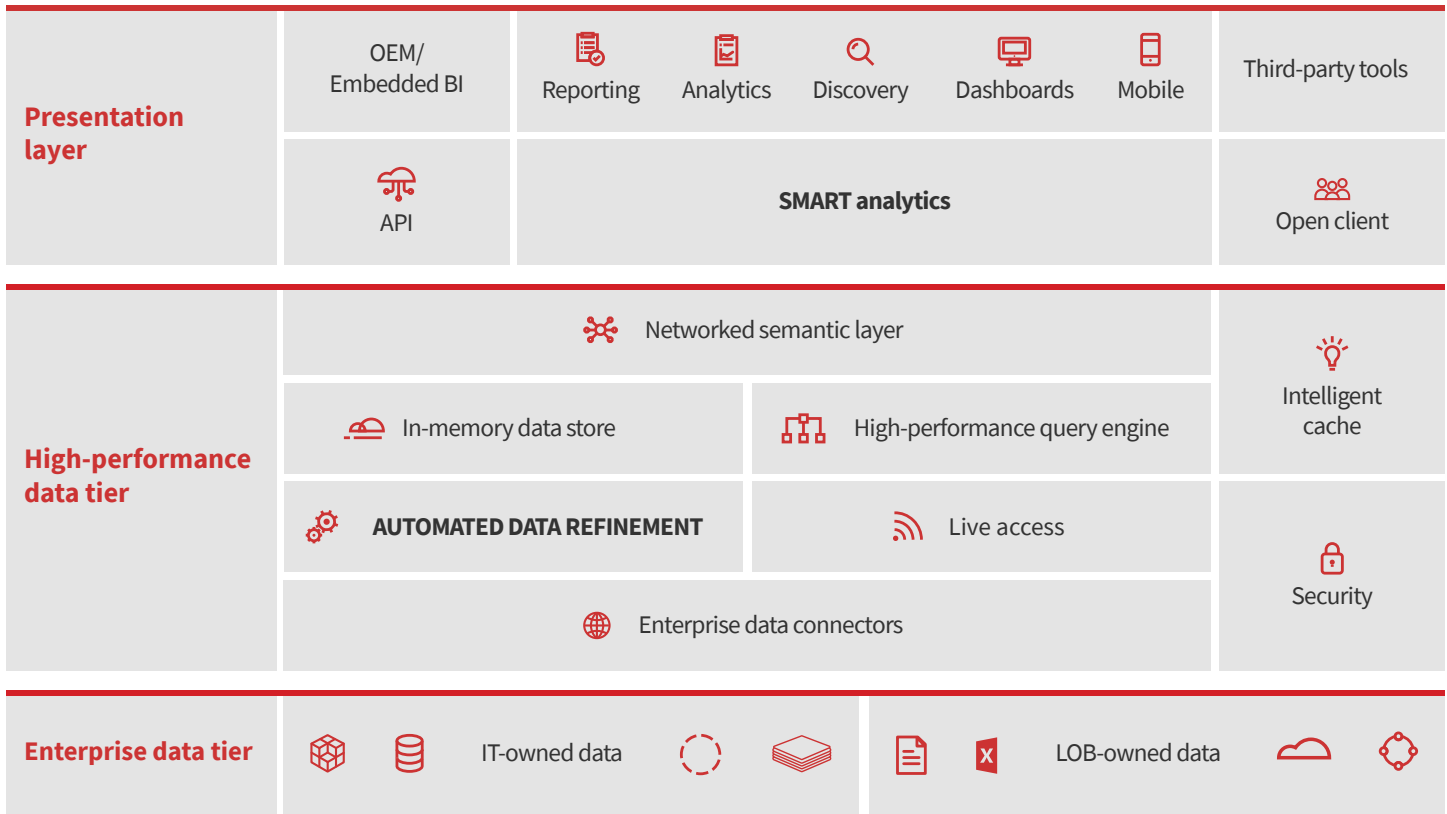
Summary

Throughout the development of the shared version of the truth, Birst's goal has been to enable a networked approach and increase developer productivity and end user flexibility. We have automated as many tasks as possible in the development, deployment and ongoing maintenance of a business analytics solution. Our patented ADR technology promotes consistency and allows for the coexistence of all types of data. As a result, Birst allows developers, administrators and end users to focus on executing more valuable tasks and less time doing BI administration.

Presentation layer

On top of the shared version of the truth, Birst provides an adaptive user experience, supporting all styles of business analytics. Birst's user experience includes visual discovery, interactive and responsive dashboards, analytic views, pixel-perfect enterprise reporting, native and offline mobile, embedded smart analytics and self-service data preparation. However, each style is not a separate tool. Birst blurs the lines between traditional analytics modules, enabling users to interact with data as they move from discovery to dashboards to reports, creating, collaborating, and publishing with a single click. Each of these styles pulls data directly from the agile networked semantic layer.

This approach enables true self-service and speed, allowing users to do more with far fewer dependencies on IT while maintaining and analyzing user-generated data. Furthermore, Birst provides an Open Client Interface that allows enterprises to utilize other analytic clients like Excel or third-party desktop discovery tools to access the Birst semantic layer, bringing governance to these desktop discovery environments.

Modern enterprise-class architecture—presentation layer

|							
Presentation layer	OEM/ Embedded BI	Reporting	Analytics	Discovery	Dashboards	Mobile	Third-party tools
	API	SMART analytics					Open client

High-performance data tier	Networked semantic layer		Intelligent cache
	In-memory data store	High-performance query engine	
	AUTOMATED DATA REFINEMENT	Live access	Security
	Enterprise data connectors		

Enterprise data tier	IT-owned data	LOB-owned data

Data discovery

Birst's data discovery interface offers the ability to intuitively explore data by creating visualizations using a drag-and-drop and double-click approach. Our intelligent recommendation engine can take the user through the process of building a visualization using the chart-first approach and providing suggestions for selecting the best visualizations. Visual filtering, user created metrics, instant metrics (like percentage of value), and intelligent search functionality add to the self- service capabilities, and the messaging center keeps users updated throughout the visualization building process.

Available visualization formats include: column, bar, line, spline, area, area spline, points, scatter, bubble, pie, donut, semi-donut, funnel, pyramid, list tables, crosstabs, heat map, tree map, waterfall, and geo maps. Advanced analytics visualizations include cluster, trendline and benchmark.

Users can apply filters and sort data to meet their needs. Data formatting is available to accommodate currency symbols, dates (including locale-based dates), decimal precision, units, percentages, conditional formatting, and on individual visual elements (axes, tool tips, and display values). The Expression Editor provides the ability to create custom expressions that can be used to create more insightful visualizations, and users can also choose to limit data to "Top N" data points.

Users can save visualizations as reports to be distributed via dashboards and embed visualizations in third-party applications. This true business-user discovery tool supports the need for non-data-savvy users to perform ad hoc analysis in a decentralized fashion while accessing the agile networked semantic layer.

Visualizer
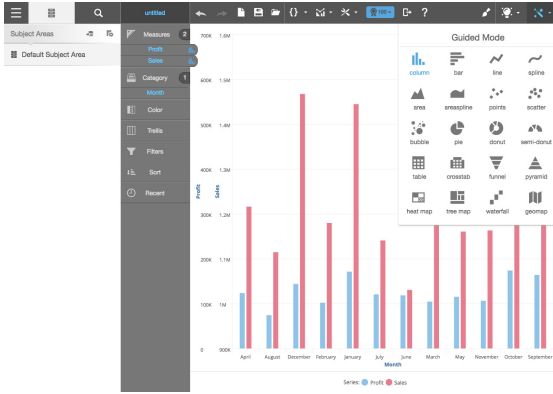
Interactive and responsive dashboards

Birst's interactive dashboards provide a self-service and easy-to-use interface for business users of all types. Unlike other dashboards, Birst provides an interface for creating dashboards that any user can leverage—with simple what you see is what you get (WYSIWYG), drag-and-drop creation—blurring the lines between dashboards and discovery. Birst dashboards and widgets are rendered in HTML5, so they automatically resize for a responsive, optimized experience wherever you use them, on your laptop, desktop, tablet, or phone. With Birst, users can directly interact with the dashboard—or even build new ones—without any formal training or specialized expertise. Even filtering has been made intuitive, incorporating filter results via prompts or lasso filters with results seamlessly cascading across filters and dashboard pages. Birst dashboards support flexible, drill-anywhere capabilities.

Both charts and tables can drill across any desired target report, dashboard or external URL. Lastly, users can incorporate external visualizations into Birst dashboards, or take any dashlet and expose it externally for embedded analytics in any web application.

Enterprise pixel-perfect reporting

Birst also includes a report designer for advanced pixel-perfect report creation, enabling highly formatted report creation typically used in production-delivered reports. Examples of rich formatting include: conditional formatting, conditional display, duplicate suppression, and null value replacement. Embedded images and sub-reports in various bands are supported. Reports are compiled into Java byte code for fast and direct execution. No interpretation at runtime is required, and server-side report caching enhances performance.

Birst reports can be exported to a variety of formats, including PDF, Excel, PowerPoint®, and CSV, with all details enacted, such as filters applied. Both business users and administrators can schedule reports for delivery by email, as attachments and in-line content. For alerts and exception reporting, you can schedule trigger reports that evaluate specific conditions. When the condition is met, for example when a KPI falls below a certain threshold, the alert email will go out. Birst also provides sophisticated report bursting, where a single database pass can be used to serve hundreds to thousands of reports, allowing high-volume report distribution without taxing the database. Birst also supports parameterization of reports based on user roles and data visibility rules.

Responsive dashboards

Pixel-perfect report

Open Client Interface

The Open Client Interface connector enables desktop- based client tools to interact with the Birst agile networked semantic layer via the Open Database Connectivity protocol (ODBC). The Birst Interface translates SQL generated by the client tool into BQL (Birst Query Language). Desktop-based analytics clients generally lack important enterprise analytics capabilities such as security and a common and reusable semantic layer and, as a result, lead to inconsistent data and information silos. The benefit of the Open Client Interface is that it enhances the capabilities of these client tools by allowing them to leverage the Birst semantic layer, ensuring a single version of truth throughout an organization. End-user experience is seamless: end users continue to interact with analytics within their client tools, while Birst executes queries in the background, ensuring every user, regardless of their client tool, gets the same trusted and governed answer to the enterprises most important KPIs.

Native and offline mobile

Birst enables users to take their insights anywhere by viewing and interacting with reports and dashboards in their mobile tablets and smartphones. Birst Mobile makes it easy to turn your data-driven insights into a compelling story. You can take your data stories on the road and deliver them anywhere, anytime. Combine interactive dashboards with custom text and annotation to create a targeted narrative designed to communicate with your audience.

Birst provides a native mobile app for Android™ and iOS tablets and smartphones that take advantage of the rich interactivity of these devices. Birst Mobile employs a "design once, use anywhere" approach.
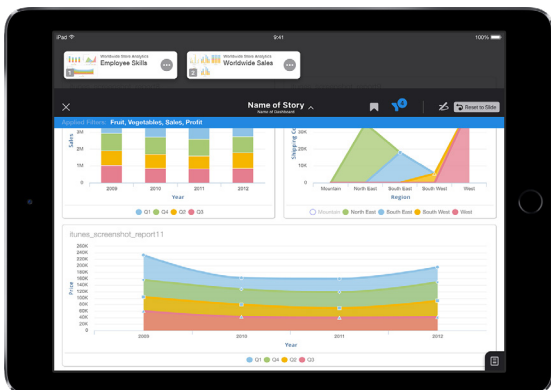
This means users develop their dashboards once, and it can be viewed on any device, either natively through an app or in a browser. Using responsive design, Birst Mobile automatically adjusts the content's layout—rotating, resizing or moving as necessary—ensuring reports and dashboards fit the device on which they are viewed. Also, Birst enables companies to white label their mobile app and use single sign-on (SSO) for a custom look- and-feel.

Birst also offers a true offline capability on mobile devices, so users can interact with their data instead of only looking at static offline images many other vendors provide. This offline capability ensures users can continue to explore their data in remote areas or in buildings where Wi-Fi is not available. To support the highest levels of security, Birst encrypts data on the device and in-transit and provides remote data-wipe capabilities.

Smart analytics

The Birst platform includes an advanced predictive analytics engine. In contrast to conventional data mining environments, data does not have to be moved; instead, datasets for model training and scoring are generated directly from the Birst agile networked semantic layer.

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

Mobile
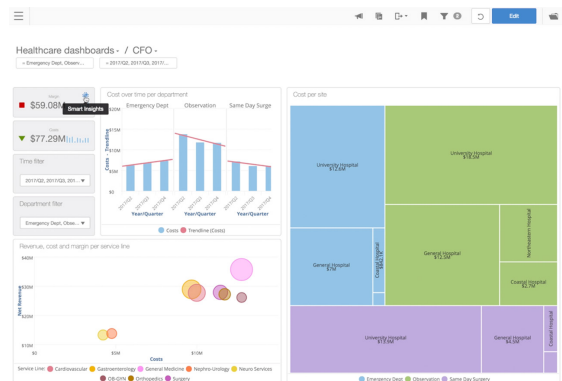
Smart analytics

Birst Smart Analytics offers a set of AI-enabled capabilities that leverage machine learning algorithms to generate automated and personalized insights for any user, regardless of their skill level. Business users can tap into the power of data science to instantly understand the factors driving their KPIs.

With Smart Analytics, Birst offers AI-driven insights that are:

- **Business-user focused**—Designed for the business user who asks the question.

- **Patented**—Built using Birst patented machine learning technology, the results are reliable from well-proven technology.

- **Personalized**—The business user selects the KPI and variables of interest, with dashboard and security filters applied to make insight generation more efficient and relevant.

- **Semantic-layer based**—Results are based on a shared or common semantic layer that is governed. Doesn't require special data processing or pre-aggregated cubes that have size limitations leading to multiple versions of the truth.

- **Networked**—Enables end-users to network their data with other data and apply smart analytics across the extended dataset.

- **Value-based design (VBD)**—Follows the VBD dashboard design approach proven to drive the greatest results, as it leads to shorter implementation times, higher user adoption and greater business impact.

Birst also delivers tight integration with the R statistics package, making it easy to operationalize R-based models to any number of users. Birst can make calls to the R server, submit data for processing, and retrieve the results to present to users. By leveraging our integration with R, users can greatly augment the already robust advanced analytics capabilities available in Birst out of the box and extend advanced machine learning to the enterprise.

Custom expressions and OLAP-style analysis

Birst enables users to create powerful custom expressions without the need to get IT involved. Birst's logical query language (BQL) allows users to both define and save OLAP-style and Excel-style calculations. This includes advanced functions, lookups, transformations, and linear regressions. For OLAP-style analysis, Birst supports aggregations, cell-based calculations, slicers, and filters. Positional calculations allow users to compare how a data point relates to values elsewhere. All analytic functions—including inheritance, business rules, multi-pass calculations, and virtual measures—are available via a point-and-click interface.

Embedded analytics

Birst empowers software providers to quickly and seamlessly embed business analytics into their applications and leverage Birst to differentiate from their competitors, delivering more value to their customers, and creating new revenue streams. Birst web services enable programmatic administration of a Birst solution and tight integration into other applications or portals.

- **White labeling**—Birst allows you to match your application's branding and look and feel on the web or mobile. Using standards such as iFrame and JavaScript, Birst provides integration capabilities to place charts and visualizations into your application. Use Birst to customize and match your application's fonts, colors, images, logos and other design elements, while still maintaining full analytic interactivity.

- **Multi-tenant**—Birst offers a multi-tenant Software-as-a-Service (SaaS) application that elastically increases in capacity as your data or user base grows. Birst also offers a fully multi-tenant virtual appliance for cases where you need to deploy analytics behind your firewall and on your hardware.

- **Web service APIs**—Birst's web services APIs extend Birst as an open platform for embedding into any SaaS or web application. Birst supports all methods of web services to receive data whether it's REST or SOAP. The outbound Birst web services API is SOAP- based and can work with any programming language that supports web services. Web services range from calls to managed users and metadata to services for running queries.

Enterprise architecture

Birst's Networked Business Analytics technology is a native cloud architecture that provides many benefits to the business, speeding time to value (TTV), reducing total cost of ownership, and increasing agility. Birst's architecture supports centralized IT teams and decentralized lines of business within large enterprises supporting the flexible demands of thousands of users and petabytes of data.

Always On

One of the key challenges with most analytics platforms is that end-users are unable to interact with data while the underlying data is being processed or loaded from other systems. Infor Birst has created a unique capability, Always On, which enables users to continuously view dashboards or visualize data while data is being processed or loaded into Birst. This ensures global organizations, or organizations frequently processing data, can interact with data and make smarter decisions without interruption, allowing for zero analytics down-time.

Global deployment capabilities

Birst is a global cloud solution, with hosting centers in North America, Europe, and Asia Pacific. Capabilities that support our global customers include multi-lingual support (translate once across multiple reports and languages), multi-currency support, multi-time zone support and multiple calendar support. These capabilities are built into the Birst platform, combined with the unique template and copy features, speed the process of rolling out Birst into new countries, by creating a single template space and extending it across multiple countries.

Birst is a fully multi-tenant solution from both a data-processing and data-storage perspective

Birst's web and application server ties are multi-tenant meaning that users are spread across an infinitely scalable pool of computing resources, leveraging its shared-nothing architecture. Birst is the only analytics solution to provide this level of scalability.

This multi-tenant infrastructure is key to enabling Birst to provide higher levels of service at a lower cost to customers, while maintaining zero cost, instant upgrades every quarter, and industry leading levels of up-time.

Networked business analytics in 78 countries

A global consumer packaged goods (CPG) company faced the challenge of rolling out a consistent approach to sell-in and sell-through analytics across 78 developing countries. While each country maintained unique data and methods of calculating different business metrics, the centralized IT team still needed to provide a single view of revenue, inventory, and demand across the different channels in each country. Leveraging Birst's shared version of the truth, the team could create a parent space, and then empowered each country to create their own child space based on their unique needs, and network it to the parent space—all while maintaining a single centralized version of key business metrics.

It also allows customers to create different logical analytics instances, all with the same physical infrastructure, drastically improving time for iterating on the development of analytics content.

Lastly, multi-tenancy also enables Birst to provide the unique virtual spaces and packages, which truly enable a decentralized user to collaborate with centralized data.

Birst's multi-tenant architecture has been demonstrated to scale linearly (i.e., a four-node configuration will support four times the workload and users) and can process larger data volumes than traditional legacy BI platforms, as illustrated in the comparison below.

Birst leverages aggregates, dynamic "indexed" cubes, and intelligent multi-tiered caching. Aggregates are automatically generated from semantic layer queries and used by the query optimizer, and they are updated automatically as part of the ETL process.

These capabilities can reduce expensive underlying database infrastructure requirements by up to 90% and therefore save development time and money compared to traditional legacy BI approaches.

- **User experience performance**—The Birst public cloud serves up hundreds of thousands of dashboard views per day. Data visualizations in these dashboards are built for end-user performance and to remove additional steps in the load process. Multiple queries are sent simultaneously from dashboards, whereas most other products send queries sequentially.

- **Data load performance**—Birst can load data daily, hourly, or even every few minutes. Birst leverages incremental loads and change detection to ensure rapid data loading and extraction. Furthermore, with Birst Always On users are still able to interact with dashboards and visualizations, while data is being loaded and processed.

- **Query optimization**—Birst utilizes push-down analytics, which moves calculations down to the database platform to leverage its capabilities, while Birst's semantic layer translates operations into database-specific functions. Also, Birst generates and optimizes queries—both for data loading and for analysis—appropriate to the backend data source.

- **Multi-tier caching**—Birst provides exact and fuzzy cache matching, as well as dynamic cube-like cache structures to help with performance. These indexed data structures provide far better reuse and generate lower database load than traditional caching approaches. The dynamic cache is partitioned amongst servers to minimize I/O contention and to allow better memory caching, ultimately resulting in a far more scalable solution. The Birst ROLAP engine in combination with this unique caching layer provides a significant performance improvement over traditional OLAP solutions.

- **In-memory MPP storage**—Birst offers multiple high-performance analytic database options, designed from the ground up for speed. This delivers massively parallel processing to coordinate processing simultaneously across separate nodes. Birst also offers support for column store databases (e.g., SQL Server 2014, Amazon Redshift) and in-memory databases (e.g., Exasol, SAP HANA) for faster and more efficient queries.
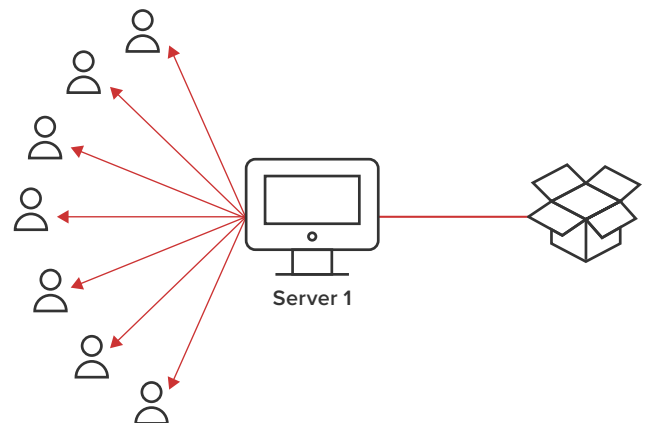
Fully multi-tenant solution

Infor Birst
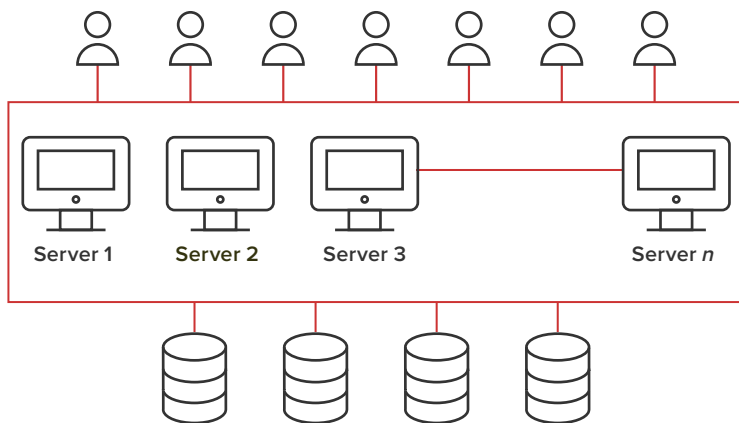
Legacy BI

- Modern web-scale architecture
- Multi-tenant environment
- Stateless application servers
- Horizontal scalability

- Client-server clustering
- Single-server dependency
- Proprietary in-memory structures
- Single point of contention (bottleneck)

Business-ready data store choices

Birst provides ultimate flexibility for customers. For those who require the fastest possible end-user performance on large data sets, they can choose EXASOL or SAP HANA as their Business-ready Data Store, ensuring that query response times on billions of rows are still sub-second. For those customers who wish to store terabytes of data economically, they can leverage Amazon Redshift and get the lowest cost/ terabyte storage fees available in the cloud. For customers with standard data sizes, Birst provides a column store analytic database that performs like a consumer web application. Lastly, companies can also deploy a hybrid model and federate queries across two different data stores. This is useful for companies that analyze data with different refresh rates. For example, historical data that is refreshed once per day can reside in a low-cost storage platform like Amazon Redshift or Snowflake, while real-time data that is refreshed every few minutes or more often can reside in an in-memory MPP platform like EXASOL or SAP HANA. Birst's ADR technology aggregates and manages this data flow automatically.

Flexible deployment options

Birst is the industry's only business analytics solution that can be deployed on-premises, in a private cloud or the Infor Birst public cloud with the same code base, upgrade path, and level of support. Users can move from one deployment model to another to meet their strategic and operational goals.

Infor Birst public cloud

Infor Birst Cloud is a multi-tenant, fully integrated SaaS solution. Users get everything required for advanced business analytics in a subscription-based package delivered in the cloud. With Birst, organizations remain agile while reaping the benefits of SaaS: fast deployment, lower costs and rapid time to value. As usage grows, Birst seamlessly expands server capacity to accommodate anywhere from tens to hundreds to thousands of users. Hosted in Tier-4 data centers, Birst Cloud requires no installation of hardware and software and is pre-configured for automatic failover and 24/7 availability and support.

Accelerators for specific use cases

Birst delivers a set of pre-packaged applications called Accelerators, which can be delivered with Birst's cloud platform and come bundled with a rich set of pre-built metadata, data transformations, measures, out-of-the-box reports, and dashboards to quickly equip end-users with a flexible business analytics solution. Birst currently supports solution accelerators for Infor CloudSuite™, sales, marketing, and operations—unifying data from Salesforce, Marketo, and Google Analytics™.

Total cost of ownership and faster time to value

Birst's cloud architecture is founded on automation, virtualization, and pre-integration, directly targeting the largest areas of cost for BI solutions. Cloud architectures completely alter how software is provisioned, configured and, deployed and offer significantly greater TCO advantages. With Birst's cloud technology, customers can deploy business analytics applications in days or weeks, not months or quarters. Birst provides for a lower TCO through reduced resources, zero hardware, zero upgrade costs, and rapid deployments. With a cloud analytics model, organizations get the benefit of SaaS—rapid time to value and upgrades in place.

Birst automates IT-centric tasks allowing customers to spend 25% to 35% more of their time on value-added activities, such as producing new reports, dashboards and rich analytics applications. Since the lion's share of costs in a business analytics deployment relates to human capital and integration costs, Birst's pre-integrated and consolidated solution drives significant TCO advantages over traditional vendors.

Security and reliability for the enterprise

Physical security

A key aspect of security is the physical security of hardware containing customer data. Birst utilizes Tier-4 data centers around the world to ensure enterprise performance, redundancy, security, disaster recovery and business continuity. In addition to making sure that the data center containing customer data is physically secure, Birst makes sure the networks and hardware containing customer data are hardened and undergo regular penetration and vulnerability testing by third parties.

Operational security

It is not enough to have a secure physical and network environment; data centers must be operated securely as well. Birst data center operational security includes policies and procedures that are SOC 2 Type 2 audited and ISO- 27001 certified. In accordance with these policies, Birst provides rich operational security across data centers and corporate processes including strict background checks and authorized-only access to confidential information, document destruction policies, change management procedures, independently reviewed disaster recovery (DR) and business continuity (BC) plans, and frequent all-employee training for information security and privacy procedures. Birst is also HIPAA attested and audited by a third party.

Application and data security

A secure infrastructure cannot protect your data if the applications providing access to your data are not secure. Birst solutions have been designed to protect the security of your information. There are two components that make up Birst application security: Authentication and authorization.

For authentication, customers can authenticate themselves to the Birst application via multiple routes: Forms-based authentication, Open ID Connect, SAML 2.0, integration with cloud portals (Salesforce, or custom single sign-on. Passwords are hashed using PBKDF2 to defend against offline attacks.

For authorization, Birst gives system administrators comprehensive security controls that can be used to control and manage the breadth of functions and features available to their end users. Birst administrators can define dashboard, report, row, and column level security to allow end users only to see the information that they are allowed to access.

Security filters allow users to share the same reports and dashboards while ensuring that each user sees only their own slice of the data. Administrators can also manage access to attributes and measures in subject areas that are controlled via user groups.

Birst encrypts all data in transit via TLS 1.1+ channels. The status of the Birst TLS 1.1+ support for all endpoints is validated daily and can be checked at any time by Qualys SSL Labs. Customer data is AES-256 encrypted at rest using self-encrypting storage.

Birst logs all login (successful and failed), logout, administrative, and database events for auditing. Furthermore, Birst automatically locks account access after a number of failed login attempts or after a prolonged period of not logging in. Birst administrators can also flag any field as auditable in the semantic layer, and Birst will log any access to that field.

Security is built into our documented software development life cycle, based on guidelines from the Open Source Web Application Security Project (**www.owasp.org**) and the SANS Institute (**www.sans.org**). Birst runs manual and automated security tests and analyzes third-party libraries for security issues on each build, utilizing third-party web application vulnerability analysis on a continuous basis (Whitehat Security), penetration testing (Whitehat Security), and third- party static security analysis (Veracode) on major releases.

Once a customer cancels their Birst account with Infor, their information will be securely maintained for the period specified in their terms of service contract. During this period, the customer can access their information only if they re-activate their account. After this period is concluded, the account data is permanently deleted from the Birst data center and is no longer accessible.

Summary

The landscape of business analytics has changed. Companies today are struggling to bridge the divide between centralized IT teams supporting enterprise requirements and user-led decentralized teams demanding greater agility. Closing this gap is the key to ensuring business analytics success.

Infor Birst's unique approach empowers business users with the speed, autonomy, and agility they demand while giving IT leaders the governance mechanisms they need to deliver a complete and consistent view of the business.

Born in the cloud, Birst's networked business analytics technology plugs into centrally managed data sources and seamlessly unifies them with data generated by decentralized teams throughout the organization. Birst then automatically refines this data and prepares it for analysis by overlaying a consistent set of business rules and definitions—creating a shared fabric of analytics across the organization—to deliver a "shared version of the truth" through an adaptive user experience that includes reporting, dashboards, visual discovery, mobile, and automatically created, personalized, smart analytics.

LEARN MORE →

Follow us :

infor™

Infor builds business software for specific industries in the cloud. With 16,500 employees and over 90,000 customers in more than 170 countries, Infor software is designed for progress. To learn more, please visit www.infor.com.

641 Avenue of the Americas, New York, NY 10011

INFDTP2155008-en-US-0320-2