The right architecture for business intelligence

The foundation for effective enterprise BI

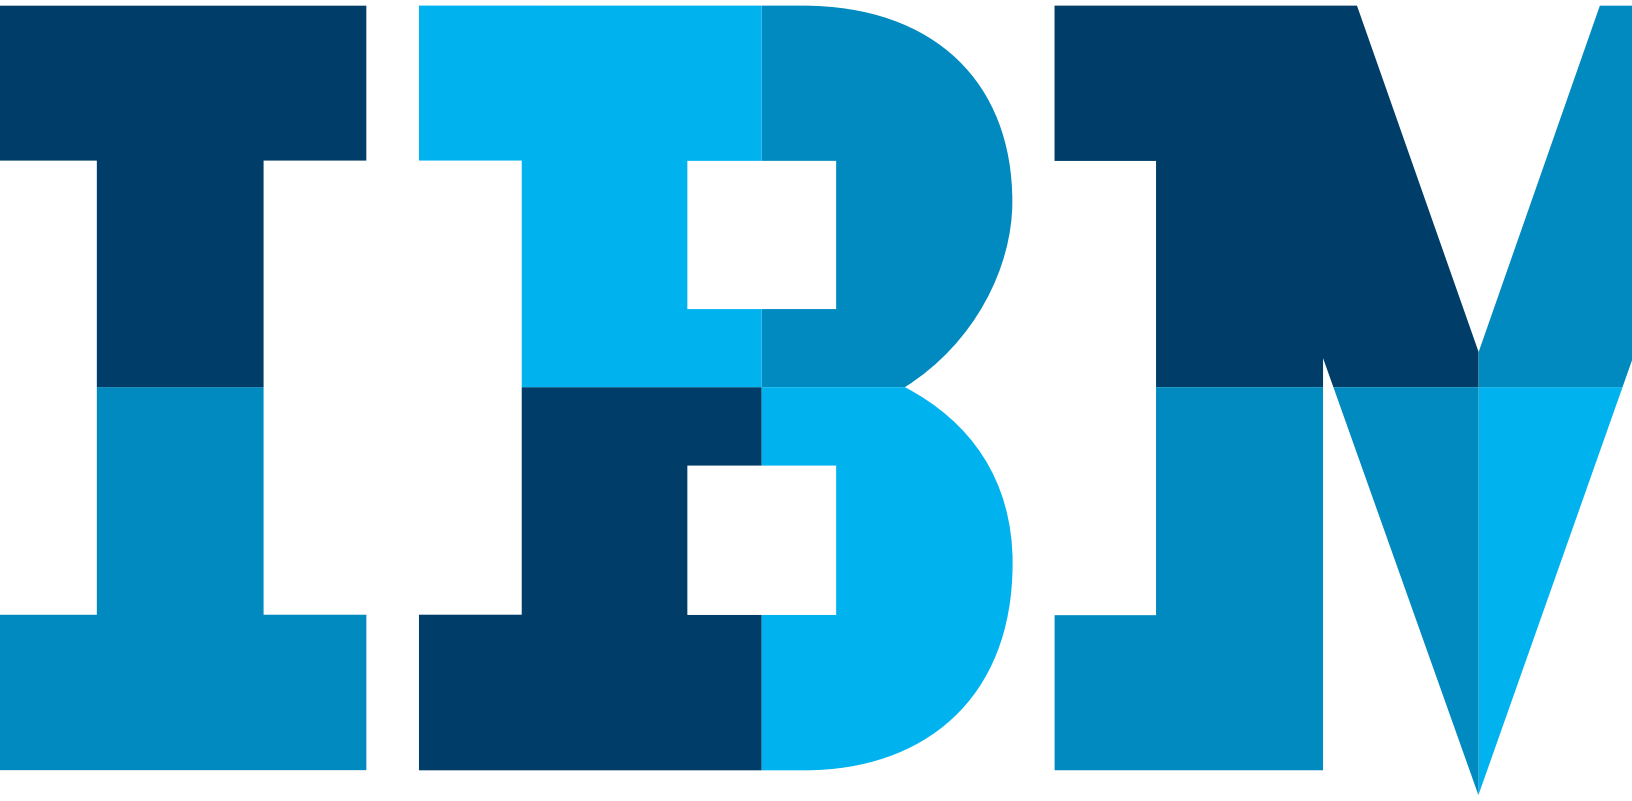

Overview

In a fast, interconnecting and complex world, it is no longer enough to decide and act on the basis of limited information or traditional strategic planning cycles. New challenges and opportunities require agility: the ability to quickly assess, reinvent and adjust.

Business analytics is helping many organizations achieve this kind of agility. Analytics software brings together business intelligence (BI) capabilities such as reporting, analysis and scorecarding with planning, scenario modeling, real-time monitoring and predictive analytics. It lets you tap into information within your organization and beyond, to connect with key stakeholders and to share insight, align and decide.

Analytics-driven organizations not only seize opportunities: they outperform. IBM's 2010 CFO study (involving more than 1,900 CFOs and senior finance leaders worldwide) showed that analytics-driven organizations had 33 percent more revenue growth and 32 percent more return on capital invested.

Investing in analytics is considered a priority for many organizations, but a big question can be the technology platform. Which platforms provide the best foundation for positive business outcomes? What kind of architecture best lends itself to accessible analysis, intuitive collaboration and actionable insight?

This white paper outlines the customer requirements that industry experts feel must be part of an effective and flexible enterprise BI architecture. And it describes the open, enterprise-class platform of IBM® Cognos® Enterprise, a solution that can change how organizations make decisions, allocate resources, predict and plan for the future, and ultimately gain competitive advantage.

Why architecture is important

A software system's architecture determines its ability to meet business needs now and in the future. The right architecture paves the way for the success of the system and, ultimately, of the organization.

An open platform based on architecture that is built specifically for analytics and designed for evolving and growing business demands meets the needs of both IT and business users.

For IT, BI software delivers more value when it

- Integrates easily with an organization's infrastructure
- Supports today's technology and standards
- Adjusts readily to evolving needs
- Consolidates all data in an organization
- Scales as user demand grows
- Performs reliably
- Provides centralized administration that does not overtax budgets and human resource budgets and human resources

For business, BI software must suit the many roles, different skills and changing needs of users. The business expects BI software to provide data in many different formats, including regular reports, responses to impromptu queries, scorecards and dashboards and yet ensure that it is "a single version of the truth" so decisions can be made confidently. In addition, the business wants BI software to be interactive, available in the office and on the go and accessible to all who need to use it.

An open platform simplifies IT environments, accelerates business decisions and provides competitive advantage by ensuring that an organization's investment can be used today and also be ready for tomorrow.

Attributes of an effective enterprise-scale BI architecture

Analyst findings and IBM's experience with Fortune 1000 organizations point to several common characteristics and values of enterprise-scale BI architecture. These requirements are fundamental to business intelligence systems that will be deployed broadly throughout the organization. All of these qualities are delivered largely through the underlying architecture.

Usability	To reach the broadest possible audience, a BI solution must provide a common user interface, recognize and accommodate different types of users, and offer all and be accessible on a range of technology, including mobile devices. It must be searchable so that users can capitalize on BI information that the organization has already created.
Seamless interoperability	There must be a single interface for all BI capabilities that can be used to navigate through scorecards, dashboards or reports. IT must be able to enable more or less functionality to fit the needs of different users.
Common business view	For organizations with many data assets, applications and users, it's critical that a BI solution deliver a common view of the business — so managers and other workers do not have to worry about the validity of their numbers when compared to other people's numbers. The single view must be based on all the data, and the quality of the data must be maintained to ensure user confidence. Data modelers must be able to create an effective business model quickly and readily modify it as the needs of business change over time.
Agility	If something in an organization changes — such as a new business strategy or a new enterprise application — the BI solution must adapt accordingly.
Scalability	Enterprise BI deployments must scale linearly to thousands and tens of thousands of users in global organization.
Reliability	For most organizations, BI is central to the business. A BI system must operate 24x7 with redundancy for all capabilities and services.
Openness	BI must be open — in terms of the data that can be accessed and for integration with existing and new applications, portals, security systems and more.
Deployability	Deploying the BI system — actually getting information to users in whatever format needed — must be a simple activity, as must making changes to the way information is deployed.
Manageability	IT must be able to administer efficiently to keep the system operating effectively. Administrators should also be able to identify and avoid potential problems.
Use existing infrastructure	A BI solution must work in existing environments and exploit everything those environments have to offer: web infrastructure, databases and OLAP data sources, security providers, application servers and more.
Security	A BI solution must work with existing security solutions — often more than one — to ensure that access to both the BI system and the information in that system is always secured as required.

Cognos Enterprise architecture

The Cognos platform underpins Cognos Business Intelligence, delivering all BI capabilities on three distinct tiers:

- A presentation tier that handles all user interaction in the web environment
- An application tier with services to handle all BI processing
- A data tier that provides access to a wide range of data sources.

The separation of the architectural components into tiers supports the secure deployment strategies required by large organizations whose data and infrastructure are secured and closely guarded by firewalls. When a request is submitted to a Cognos Enterprise installation, the right processing happens at the right level and in an optimal way to serve the broad range of business users. Processing on the presentation tier, for example, can mesh with existing load balancing routers, ensuring that as requests come in they are distributed appropriately.

The Cognos platform is built with web services to deliver BI from a single extensible and flexible platform.

Figure 1: IBM Cognos Enterprise: all BI capabilities on a common modern architecture.

The presentation tier

The Cognos platform provides all BI capabilities in a web browser-based user interface and makes other capabilities available in a self-service, standalone desktop interface. BI users, BI authors and BI administrators can use their web-based workspaces or desktop dashboards to read reports, to create ad hoc queries and analyze data.

Unlike many other BI tools, there are no applets to download and no plug-ins to install or maintain. Instead, simple web gestures build and access a broad range of BI content, which helps improve the productivity of many types of users in your organization—from report authors to business managers and your front line. In addition, Cognos consumer modes use the Cognos platform to provide business users access to mission-critical BI on mobile devices and in familiar software applications such as Microsoft Office.

The administration of Cognos Enterprise is also browser-based, which reduces the administrative burden on IT. IT does not have to install and manage client desktop software, thereby minimizing deployment and maintenance costs. Administrators can manage and tune servers, add groups and users and grant privileges to secured business intelligence content. For large organizations, which are often geographically dispersed, this means administration can be distributed so local or regional administrators can handle updates for users and roles while others manage overall security centrally.

The Cognos administration interface provides IT with views of BI system activity so they can manage the system proactively. Intuitive, at-a-glance metrics, role-based capabilities, queue prioritization and resubmission features help ensure that IT can keep their BI system functioning optimally.

Infrastructure components	Cognos Business Intelligence benefit
Web server	Use any popular Web server including Microsoft IIS, IBM HTTP server or Apache. The Cognos platform fits into these environments with lightweight web gateways that forward incoming requests from the web browser to the application tier.
Application server	Exploit the power of in-place application servers, including BEA WebLogic, IBM WebSphere® family, SAP NetWeaver Application Server or Oracle Application Server.
Routers	Mesh cleanly with in-place load balancing mechanisms for optimal use of hardware resources.
Portals	Integrate with widely used portal environments such as IBM WebSphere, Microsoft SharePoint, SAP NetWeaver and BEA Plumtree. In fact, the Cognos platform was built to integrate with portal environments that comply with the Web Services for Remote Portlets (WSRP) standard. As a result, organizations can integrate the Cognos platform into their enterprise portals — today and in the future.

Reducing costs by using existing web infrastructure

Almost every organization has a web infrastructure, and Cognos Enterprise must fit with that infrastructure. By working in whatever web environment is currently in place—and thereby using existing skills and assets—the Cognos platform helps reduce the effort and costs associated with getting an enterprise solution up and running.

Integrating business intelligence with existing applications

Although analytics has become truly strategic in many organizations, it must fit in a framework of existing business applications. With a fully open and documented application programming interface (API), the Cognos platform provides a significant level of openness. You can integrate most BI capabilities into your existing system and choose from widely used programming languages such as Java, C++ and others.

The Cognos Software Development Kit (SDK) exposes the same web services API used to build Cognos Enterprise. You can also access the API with WSDL, and any programming language that understands SOAP can use it. The BI content in a Cognos platform configuration can be integrated with Java-based JSP applications or in the Microsoft .Net framework.

The application tier

The application tier is the mission control center of the Cognos platform. It manages all incoming interactive and batch requests and automatically distributes them for optimal impact. It also provides a single set of standard-based services—such as a common query engine, scheduling, monitoring, auditing and presentation.

Self-registering, self-starting servers

When you are configuring an enterprise-scale system, you must make sure that it can maintain the best possible level of service. For this to happen, incoming requests should automatically go to the appropriate server for best throughput.

In the Cognos platform, the optimal routing of requests in the application tier is the job of the dispatcher. The dispatcher is a multithreaded application that runs on whatever web application server or servlet container an organization is using. These include Apache, BEA WebLogic, IBM WebSphere Application Server, SAP NetWeaver Application Server and Oracle Application Server.

Each dispatcher in a distributed system is self-registering. When you install the Cognos platform on a server, the dispatcher simply registers itself in that configuration, starts the services on that server and lets the configuration know what services are available. This vastly simplifies the installation and configuration of Cognos Enterprise and it enables the system to scale even if it is installed on multiple servers.

Intelligent, configurable load balancing

Enterprise-scale BI systems must be able to handle the high volumes of incoming user requests typical in large organizations. Requests to run a report, to display a dashboard, to burst a scheduled report to a wide number of users and more should be answered with optimal performance.

In the Cognos platform, requests are dispatched with load balancing built into the system. As requests come in, they are automatically routed to servers based on defined server capacity and the request's "affinity" level, which the dispatcher uses to decide whether the request should go to a specific server or to any server in the configuration. This affinity can be derived from the actual nature of the request, or from the group or user role of the individual submitting the request. As a result, servers can be dedicated to specific groups or users.

The capacity definition for any given server is completely flexible—if one server has twice the "power" of another in terms of memory and processing speed, it will automatically have twice as many requests dispatched to it. Additionally, every server in a Cognos configuration can be tuned to adjust specific performance parameters such as the number of active request threads for any given service, timeouts and the level of auditing applied to any given BI activity.

Peer-to-peer services for reliability and scalability

The services offered by the Cognos platform are the backbone of the system. Regardless of what kind of request is made—a simple report run, an analytical comparison across business dimensions or the scheduled running of a business intelligence agent that detects key data events—the system must provide services smoothly.

Every service in the Cognos platform application tier operates on a peer-to-peer basis. This means that no service needs to know any of the details about what any other service does or is doing at any given point in time. Any service, on any machine can handle any incoming request. It also means linear performance characteristics, unlike other SOAs that use a services "hub."

The nature of these services is such that there is complete separation of elements that should not be tightly bound, such as presentation and data. Presentation is handled by a presentation service, and data is handled by the query service, based on the built-in business rules in metadata and in defined security.

The result is fault tolerance and service redundancy—any request can be routed to and handled by any server in the system. If any server in a configuration fails, incoming requests are automatically routed to redundant servers, thereby avoiding service interruptions. The services are also scalable, so administrators can add servers and enable or disable services based on demand. For example, it's a simple matter to dedicate a specific server in a Cognos Enterprise configuration to report execution by disabling the other services on that particular server.

BI Bus

The open API for integrating the Cognos platform into other systems is used by all of the components and services. All the communication between the services in the Cognos platform configuration takes place on the BI Bus—which means that all services plug into a "network."

As a result, services are transparent in terms of location. Services communicate with one another with common messaging that uses open web standards: HTTP, SOAP, XML and WSDL. The intra-service communication is coarse-grained in nature. This means each request typically handles a significant block of work, and intra-service communications is optional. Additionally, intra-service calls can be encrypted, ensuring security in the application tier.

Single authoring model and common metadata for consistent results

Having a single authoring model based on common metadata, regardless of where or how that data is stored, is fundamental to a successful enterprise BI solution. If a solution has no common understanding of the data and uses multiple authoring models—for example, one access mechanism for production reporting, a second access mechanism for multidimensional reporting and possibly a third for ad hoc query capabilities—then the very real possibility exists for inconsistencies in these various BI activities.

The Cognos platform uses a single authoring model for both relational and dimensional data source that is based on common metadata. Users can therefore have confidence that the numbers in their reports will match those from other departments. Regardless of whether a user is accessing a relational data warehouse or a multidimensional data cube, the query engine will use defined metadata and generate underlying queries that return consistent results.

High performance, in-memory caching with dynamic query

Organizations are challenged with meeting performance expectations because of complex business requirements and growing data volumes. Dynamic query provides in-memory processing capabilities with 64-bit system support to address these challenges. This capability brings optimized query generation with pattern intelligence and security-aware in addition to providing the foundation for further extension such as dynamic cubes.

Dynamic query uses an enhanced Java-based interface that addresses query complexity, data volumes and performance with new capabilities:

- In-memory calculations and aggregate operations
- Smart query processing capable of combining multi-dimensional and relational concepts to improve performance
- 64-bit security-aware smart caching that uses in-memory optimization to increase query performance and data cache reuse and that monitors the most commonly used data and responds accordingly.

Cache management facilities are available by using existing event scheduling infrastructure to enable the automatic management of the cache, so the content can remain relevant.

High-performance analytics over terabytes of warehoused data with dynamic cubes

As data volumes grow, so do performance challenges. Users expect very rapid responses, especially for dimensional interactive analysis. Dynamic cubes, an extension of dynamic query, address these performance problems by means of aggregate awareness and extensive in-memory acceleration.

Aggregate-awareness enables you to fully capitalize on the investment made in a data warehouse technology, where different types of data are brought together in a structured relational schema. Query routing is automated, so that authors and analysts can simply benefit from enhanced performance without worrying about whether their query will hit the right database tables.

In addition to aggregate-awareness, in-memory caching further accelerates performance, giving IT the tools to build a flexible and optimized application.

In-memory aggregates can be created with the Aggregate Advisor and no re-authoring nor re-modeling is required. And a sophisticated caching strategy uses metadata, individual data points, expressions and result sets to optimize performance.

As with the rest of the platform, this technology provides you with administration tools to manage environment variables and observe system metrics on a granular level, along with scheduling options and full SDK openness to automate management tasks such as cache priming or cache refreshes.

Platform independence for flexibility

Another key element of the Cognos platform is environment independence. In terms of operating systems, you can use your existing infrastructure and install the Cognos platform on Microsoft Windows, UNIX, Linux or IBM z/OS®.

If you have multiple operating systems—for example Windows and Linux—you can configure your BI system for a heterogeneous environment.

The data tier

Large organizations typically have multiple data sources. At a departmental level, there might be a huge proliferation of data sources that makes the delivery of business intelligence on an enterprise scale difficult. Most organizations have both relational data and multidimensional data. They might already have significant investments in metadata. Even organizations that have managed to standardize their data strategy are potentially subject to multiple data sources as soon as they merge with another organization or choose to grow through acquisition.

Many BI systems provide access to some of these sources. But IBM offers access to many more of them. As a result, you can:

- Use all of the data sources at your disposal.
- Work with federated views of your multiple data sources.
- Capitalize on existing enterprise data warehouses or data marts, with the ability to create new ones.

Centralized, maintainable, secured BI content

The assets managed by a BI application are critical to an organization's business infrastructure—just as important as the underlying data assets in ERP systems, in relational databases and in other data sources such as XML streams or web services. As with other critical assets that are used to manage the business, a loss of BI content can be very damaging.

In the Cognos platform, all BI content is stored and maintained in one location: the content store. As with virtually all critical information assets, the best place to store BI—including reports, metadata packages, configuration information, user and group preferences and key metrics—is in a relational database management system. All of the value of a relational system, which includes performance tuning, security, backup and recovery and global accessibility, can then be brought to bear on BI applications.

Cognos Enterprise includes a content manager cache service that enhances the overall system performance and Content Manager scalability by caching frequent query results in each dispatcher. This service can be customized.

The Cognos platform employs widely used relational databases as the storage mechanism for all BI content. Depending on your needs, BI content can be stored in IBM DB2® in Oracle, in Microsoft SQL Server, in Derby or in Sybase. And, as with all BI services in a Cognos platform configuration, redundancy is built into the system, with multiple instances of the content store for failover and reliability.

Common metadata for a common view of the business

The Cognos platform provides a common metadata view. Metadata modeling capabilities help IT groups build enterprise-scale metadata models to meet your requirements. IT can use Cognos Framework Manager to build data models that span a wide range of data assets and provide a common view of data. A single metadata model can be built on metadata derived from diverse data sources such as Oracle, Microsoft SQL Server, Erwin and XML. Teams of modelers can work independently on different parts of a model and combine their work. They can also use a single model to send different packages of information to different types of users.

It's important to note that when importing metadata from various sources, Cognos Framework Manager uses everything it can from the data source in terms of metadata, including joins, cardinality, dimensions, hierarchies, attributes and measures. For example, when you import from IBM DB2 Cube Views, virtually all of the dimensional information inherent in that system is brought into Cognos Framework Manager which speeds the journey to metadata-driven BI.

Multilingual capabilities and UNICODE for global deployments

A core design principle of the Cognos platform is global deployability.

Support for global deployments is built into the metadata layer. With no coding whatsoever, metadata models can drive multilingual deployments. Reports, in-depth analysis, dashboards and scorecards can be delivered in any language or locale from one UNICODE server. At runtime, the local settings in a user's browser direct Cognos Enterprise to render results in the appropriate language, using appropriate locale settings for variables such as currency and the formatting of monetary values.

Use in-place security assets

Every BI application of any scale must be secured. Regardless of what form the data takes, it must be seen only by those authorized to do so. Moreover, where the security of information over the web is a concern, information must be encrypted to a level that assures the organization that its data assets won't fall into the wrong hands.

The Cognos platform uses a wide range of in-place security assets, in three critical areas:

- *Authentication.* The Cognos platform uses whatever authentication mechanism or provider is in place, regardless of how users in your organization log on to the system. It supports Microsoft Active Directory, Windows NTLM, Netegrity SiteMinder, LDAP, existing Cognos namespaces or combinations of these where multiple security providers are in play.
- *Authorization.* In the Cognos platform, security can be applied at virtually any level, starting with secured access in the metadata model to query subjects, to rows, to columns or to entire published business intelligence packages. Additionally, in the common portal environment, security can be applied to specific objects and capabilities, such as reports, analysis, dashboards and scorecards and agents. It can also be applied to folders that contain any combination of these objects.
- *Encryption.* In many BI environments, encryption is a basic requirement. The Cognos platform provides cryptographic services that apply to all information, including transient communications between services and static or temporary data artifacts generated by the system. The standard cryptographic provider employs SSL and includes trusted communications with digital signing of SOAP-based messages on the BI Bus. Strong encryption of up to 1610 bits is available from cryptographic providers.

An important distinction with the Cognos platform cryptographic services is the "across-the-board" encryption capability. If required, all inter-service communication between services in the application tier can be fully encrypted.

Comprehensive auditing

Logging is fundamental to many BI applications, for example, to meet SOX reporting requirements and for audits. Charge-back requirements often mean you need to know who is hitting what system, when and for how long.

In the Cognos platform, comprehensive auditing is provided for all services, and audit results can be centralized. All logging from all servers can be directed to one location if desired. Auditing levels are adjustable and can be set or directed to a location of your choice, such as databases, UNIX System Log or Windows Event Viewer. The auditing model and sample associated reports are based on published schema and provided out of the box.

Conclusion

Analytics support how people think and work, giving them the ability to find the right information, gain insight, share it with others and see the business from any perspective. To be an effective foundation for analytics, the right business intelligence architecture must support this intuitive and effortless interaction with information.

Guided by an open data strategy and backed by an industry leader, the open, enterprise-class IBM Cognos platform provides a strong foundation to easily deploy, use and integrate Cognos Enterprise successfully. Your organization gains the freedom to see more, do more—and make the smart decisions that drive better business results.

About IBM Business Analytics

IBM Business Analytics software delivers data-driven insights that help organizations work smarter and outperform their peers. This comprehensive portfolio includes solutions for business intelligence, predictive analytics and decision management, performance management, and risk management.

Business Analytics solutions enable companies to identify and visualize trends and patterns in areas, such as customer analytics, that can have a profound effect on business performance. They can compare scenarios, anticipate potential threats and opportunities, better plan, budget and forecast resources, balance risks against expected returns and work to meet regulatory requirements. By making analytics widely available, organizations can align tactical and strategic decision-making to achieve business goals. For further information please visit **ibm.com**/business-analytics.

Request a call

To request a call or to ask a question, go to **ibm.com**/business-analytics/contactus. An IBM representative will respond to your inquiry within two business days.

IBM Canada Ltd.
3755 Riverside Drive
Ottawa ON K1G 4K9
Canada

Produced in Canada
December 2012

Please Recycle

YTW03092-CAEN-01