


IBM i

Intelligent Business for the IBM i Enterprise

A 'need to know' guide on selecting the right B.I. solution
for the modern business running on IBM i

Andrew Nicholson, April 2013



Intelligent Business for the IBM i Enterprise

A 'need to know' guide on selecting the right B.I. solution for the modern IBM i business

Introduction

The IBM i has long been thought of as the platform of choice for the savvy enterprise. This fully integrated server technology offers some of the greatest resilience and transactional processing performance with the lowest total cost of ownership available.

However, while our IBM i enterprises enjoy class leading reliability and security, we fall short when it comes to extracting and analysing information to better enable decision making.

This has not always been the case as since the beginning of the AS/400, users enjoyed Query/400, an easy to use querying solution that was met by envious eyes from the Microsoft community in the late 1980s. But that was then and other platforms have since caught, and overtaken Query with native solutions from their own and third party stables.

In 2007, IBM purchased one of these third parties, Cognos and now had a match for high-end windows based B.I. In 2008, through the introduction of DB2 Web Query, IBM addressed their native IBM i shortcomings too. But its success so far has fallen short of its AS/400 cousin and Web Query was labelled by members of the community as far too complex and too costly.

Other up and coming B.I. vendors were quick to throw their hat in the ring too but the majority of which were windows based solutions too; but does this matter and what are the benefits of having a native IBM i solution anyway?

If there are readily available tools for SQL, then shouldn't we just port all the information out of the IBM i, for our business intelligence requirements?

What of these tools will I need, how can I view this information and what is B.I. *really* anyway?

This document will tackle these questions and suggest what a business on the IBM i should look for in a modern Business Intelligence solution.

The Competitive Edge

In these tough economic times, businesses need all the tools available to retain and enhance their competitive edge. Their ability to identify trends, predict results and understand potential is the difference between a successful and failed business.

Business Intelligence (B.I.) solutions provide these tools that better enable businesses to compete and prosper in difficult operating environments. Not having pertinent information to hand or reacting slow to insight could have a real impact on the bottom line through loss of repeat business and new sales. B.I. tools provide measurement and statistics and allow for better decisions to be made; decisions through fact-based analysis rather than second hand data or 'gut feel'.

Indeed, there is the added risk is that your competitors are making these kinds of investments and are able to react more quickly, manage their business better or respond to changes in the market better and, as a result, perhaps have better earnings than you.

As such, B.I. should be central to your company's growth strategy. The very nature of adoption could provide the difference between stagnation and achievement.

A Measurement of Success

So what is involved with business intelligence anyway and how do we use these tools to measure business actions such as customer satisfaction? These days, B.I. comes in many forms but in the main, is essentially used to the provision of clear analysis of the company's performance through reviewing different data results (financial, manufacturing, warehouse, etc.) against key performance indicators.

Role based dashboards can be designed to show live gauges, graphs and diagrams to view the information you want, immediately. From here, you can drill down to different levels of data to get the complete picture.

Information into Insight

With almost all business running on IBM i having a database full of data, getting information from its raw transactional form into insight remains a stumbling block. While query can provide simple reports, the key to sustained profitability is gained from true business intelligence; meaningful insight and foresight from multi-dimensional analysis in areas such as key performance indicators, financial projections and demand prediction.

It's more than the end form however. Each business stakeholder from executives and staff, partners and suppliers, should have access to pertinent

and meaningful information in a format of choice, when, where and how they want it.

This ability to access information when and how, plus the ability to gain insight and analytics provide the user, supplier, partner and business alike with the ammunition to collaborate, develop strategy and drive change for higher profits.

Opening doors to Windows

Let's get this out of the way first; Windows has more business application written for the platform than IBM i. The simple truth is that if you're looking for choice from big brand vendors, then the majority of big brand vendors are written over and for the Windows platform.

The Native Need

For 25 years, IBM i has paved the way for integrated architecture with its unified database, operating system and server. As such, it enjoys many idiosyncrasies that are unique to platform.

The truth? You can't handle the truth!

A common misconception is that an enterprise running on IBM i MUST create a data warehouse. But why? And what exactly is a data warehouse anyway?

A data warehouse is a central repository of data which is created by populating data that's required for reporting from the enterprise system (or other disparate sources) to an external, separate database which will be used for analytic purposes only.

As the mainstay of B.I. solutions are developed over a Microsoft platform and require an SQL database, routines to 'port' over the relevant data from the IBM i enterprise to the data repository are required. These ETL routines (extract, transform and load), extracts and organises data from the source database then transforms the extracted data to map, join and sort various data set before validating and loading in a structured order onto the SQL warehouse.

These routines are normally set-up to run once a day / week /month, depending upon enquiry requirements.

As such, the end user is constantly looking at historic data and there will always be a discrepancy between what the user is analysing and the enterprise truth.

The IBM i and BI technology conundrum

The plethora of different technologies is

Don't just query, enquire!

Querying is a great method to ask a specific question on a specific area. It allows us to set-up standard questions that an enterprise may wish to ask on a daily/weekly basis. However, should the enterprise wish to expand on these questions and try and get more meaningful data from the questions they ask, then a different method should be applied.



Underneath the TurboCube Technology



ShowMe, the popular browser based business intelligence solution for the IBM i, incorporates a feature called 'TurboCube'. This sophisticated technology was built in-house by the Utilities 400 development team and is designed to leverage the power of the IBM i platform. It employs a sophisticated engine that allows the user to refresh a business cube without the need to rebuild. As such, this technology enjoys the double benefit of drastically lowering server load while simultaneously reducing the time required to interrogate data. But how does this technology work? askBOB caught up with Andrew Nicholson and Steve Close of Utilities 400 to learn more.

Steve explained, "ShowMe builds cubes of data on the iSeries from the transactional databases. These cubes reside in a physical file on the iSeries and can be either fixed format (The designer defines all the valid routes through the data) or variable format (The user can start anywhere and drill anywhere)."

Andrew chipped in, "This allows us, in a true sense of the word to build 'one version of the truth'. Although we can build data warehouses, this technology can negate that path should that be the preferred method for the customer."

Steve continued, "Using our turbo cube technology it is only necessary to populate the cube once and then ShowMe will just apply net changes to the cube. It does this by processing journal receivers for the files that are used to populate the cube. The designer can decide how often these journal entries are to be processed - daily, hourly or even *IMMED which means that as transactions occur they are reflected in the cube within seconds. This is achieved by subtracting the before image of the journal entry from the cube and then adding in the after entry."

"If the files are currently being journalled, for example for a High Availability solution, then ShowMe will read the existing journals on the system. If the files are not journalled then ShowMe will completely manage the journalling process, starting the journalling, changing the receivers when necessary and deleting the receivers when they are no longer needed by ShowMe."

"Using this methodology removes the need to continually refresh the cubes. In fact the only time a refresh would be needed is in the instance for some of the entities being re-defined and the user wishing to rebuild the historical data to reflect the new structure."

Andrew concluded, "It's about providing the business the ability to enjoy their data in real-time, not sometime. TurboCube not only accomplishes this, but reduces processor usage. As such, compared with other big named Business Intelligence solutions, this may save over 1,000% in build and interrogation time while simultaneously lowering server load."

So there we go, in a true sense – an IBM i based business can really herald significant advantages over their Wintel or Unix based competition. A real 'i' opener indeed!