
Dynamic Business Intelligence

Leveraging mainframe assets to help steer the business

Tony Lock and Dale Vile, Freeform Dynamics Ltd, March 2009

One of the most versatile assets many organisations have in their IT infrastructure is the mainframe, which in recent years has evolved to deal effectively with a wide range of workloads. One workload that mainframe users may wish to consider again is data warehousing. In this paper we look specifically at how the latest IBM System z capabilities could allow the mainframe to play a more central role in a coordinated business intelligence strategy. We consider factors such as the potential use of the mainframe in managing operational risk and control, the use of real-time BI, and the increasing focus on the optimisation of running costs.

KEY POINTS

Process latency in the delivery of business information is no longer acceptable

Gone are the days when organisations could rely on a handful of analysts to periodically process 'the numbers' on business performance to generate monthly management reports. In today's business environments, relying solely on this approach cannot maximise efficiency or competitiveness, as those driving the business need to know about important developments before it is too late to act. The need nowadays is for an on-demand, holistic, real-time view of business performance. As a consequence, the modern business intelligence (BI) infrastructure must deliver consistent information, proactively across the organisation in a timely, secure and cost effective fashion.

Flatter organisations, empowered workforces and more accessible technologies have conspired to create new challenges

The flattening of organisational structures and the devolution of decision making tasks means that more people today require management information to do their jobs. However most organisations are not geared up to cope with this 'democratisation of business intelligence' – their core infrastructures are not designed to deal with the rapid, secure and broad information delivery. Numerous 'small footprint' solutions have sprung up independently at departmental, workgroup and even individual employee level, causing significant information fragmentation. This leads to manual overheads, incomplete views of the business, multiple versions of the truth, user frustration and ineffective processes.

A more coordinated architectural approach is required, but it's not that easy

The time has come for many organisations to adopt a coordinated architectural approach to business intelligence. While deploying yet another local, independent data warehouse may appear a good idea, this risks perpetuating the issue of fragmentation. A more centrally managed approach ensures the enterprise operates using a secure, cost effective, single version of the truth which can be built upon, rather than every system creating its own story and requiring individual management.

The answer for many might already be sitting in their data centre

For those organisations that already have a mainframe occupying the data centre, a pertinent question is whether it could provide the joined up capabilities that may be lacking. Recent advances in the IBM System z architecture, coupled with associated developments in software and tooling, plus significant changes in mainframe pricing options can now provide a linchpin to enable a more coordinated, secure, and cost effective architectural approach to solving the business intelligence challenge, especially in scenarios where much of the business data resides on the mainframe, thereby avoiding many costs and risks associated with data movement.

This paper was commissioned by IBM. It was authored on an independent basis by Freeform Dynamics drawing upon relevant findings from a number of primary research studies, as well as the extensive experience of the Freeform analyst team.




Contents

Contents.....	2
Introduction	3
Challenges and imperatives in more detail.....	3
Information, information everywhere!	3
Pleasing all of the people all of the time	4
Tell me what's driving the business now, not what drove it last quarter.....	4
Just tell me what I need to know, when I need to know it	5
And by the way, remember we are in an economic downturn.....	5
Infrastructure evolution.....	6
Possible approaches	6
Key platform infrastructure considerations and requirements	6
Upping the level of currency and responsiveness	6
Dealing with the access requirements of a broader audience.....	7
Data warehouse performance, scalability, resilience and security.....	7
BI / Data Warehouse Evolution.....	8
Mainframe Architectural options.....	8
The architectural development of the Mainframe as an Enterprise Platform	9
The Mainframe, Data warehousing and BI	9
System z-centric BI / DW	10
BI / DW on Hybrid System z Architectures	11
BI / DW not utilising System z.....	11
Relevance of options.....	12
Where is a mainframe centric approach to BI / Data Warehousing appropriate?	12
Where is a System z / hybrid approach to BI / Data Warehousing appropriate?	12
Where the mainframe approach to BI / Data Warehousing does not fit?	13
Recommended next steps.....	13
The role of Financing.....	14
Discussion.....	14
References	15
About Freeform Dynamics.....	16
About IBM	16
Terms of Use.....	16

Introduction

Over the years, most large organisations have spent significant sums of money on business intelligence solutions, from huge centralised data warehouses (DW) to a range of more distributed capabilities based on both commodity components and specialist tools. Yet despite this investment, there are still widespread complaints from business managers and other users that require business performance information to do their jobs. Even those industries best known for their advanced use of IT, such as financial services, have to grapple with some serious challenges in this area, illustrated by still-pertinent feedback from senior decision makers in the City of London¹¹ (Figure 1).



This picture demonstrates a state of affairs that is not particularly impressive, but is indicative of a general problem that touches the majority of large organisations.

As a starting point, this paper considers the nature of these challenges and how they have arisen. It goes on to provide advice and guidance for their resolution, based on a more coordinated architectural approach to addressing business information requirements in a robust, secure and future proof manner. We shall be paying particular attention to the potential role that an important existing asset could play, which many organisations already have in their armoury – namely the mainframe. Whether this asset has a place in this context is by no means a given, but with recent developments which have added many BI / DW capabilities, it is an important area to explore.

Before getting into this, let's look at some of the practical considerations in more detail.

Challenges and imperatives in more detail

To appreciate some of the later discussion on solutions and architectures, it is useful to have an understanding of the issues and imperatives involved. The major challenges faced by organisations in the areas of BI and data warehousing include the fact that data, and hence information, are now stored in multiple locations making it difficult to obtain a consistent, accurate picture of what is going on. This is a matter of great concern given that BI can no longer be considered only as a backwards looking review of operations; indeed, it is deployed so organisations can make forward looking business decisions as a routine part of everyday operations. Furthermore, the increasingly critical nature of business decisions made using BI means the security and accessibility of information-generating systems is now a matter of great importance.

Information, information everywhere!

At the heart of many of the problems discussed above is the tactical approach taken over the years to address BI requirements to serve the local needs of individual departments, workgroups and even single users. Fuelled by the challenges associated with 'boil the ocean' corporate data warehousing initiatives made popular in the 90's, and coupled with the availability of relatively cheap and accessible small-footprint technologies that provided an alternative 'fast track' route to delivery, we have witnessed the uncoordinated proliferation of data warehouses and marts across many organisations.

And the end result? Apart from the obvious problem of the wheel being reinvented repeatedly from a development and integration perspective, the business problem can be summed up in one simple phrase – ‘information fragmentation’.

We can see this clearly from the issues highlighted in Figure 1. Fragmentation makes it difficult for those who need information to find it, and find it quickly enough to make a difference. Often, people know the data they require is there somewhere, but locating it is a real issue. There are then the time delays as people in various parts of the business manually collate information for analysis from disparate systems. Lastly, there is the problem of data inconsistency, which so often leads to “multiple versions of the truth” being used to support business management and decision making processes; a very risky and undesirable state of affairs.

Pleasing all of the people all of the time

There is a quote attributed to Abraham Lincoln, “You can fool some of the people all of the time, all of the people some of the time, but you can’t fool all of the people all of the time”. But tell that to someone trying to maximise performance against aggressive targets when they can’t get hold of the information they need. The fact is that today, with the flattening of organisational structures and both responsibility and empowerment being pushed deeper into and more broadly across and outside the business, the audience that needs to be served today by BI capabilities is far greater than before.

A major consideration is the nature of the user. Unlike in the past when business intelligence was about feeding data to a relatively small team of business analysts, today we talk very much about the ‘democratisation of business intelligence’. What we mean is addressing the need and expectation of the broader community of people participating in management and decision making activity, to access directly the information they need when they need it.

It also goes without saying that such individuals want to access information via familiar mechanisms such as their desktop office tools or their browser, as they can’t (or indeed, won’t) spend time learning how to use complex specialist analytical tools. Further, such democratisation is not restricted to within the enterprise itself, as today it is often the case that customers and partners now routinely demand access to information which until recently was confined solely to internal constituents.

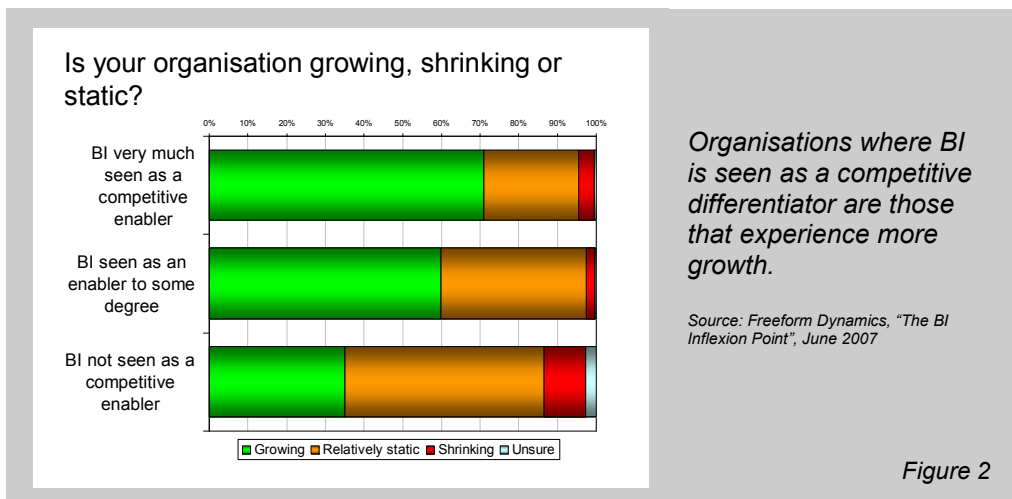
Tell me what’s driving the business now, not what drove it last quarter

The time factor has been mentioned a number of times so far, which is apt because this is one of the most prominent complaints from users of business intelligence solutions. To understand what is behind this, we must consider the concept of cause and effect. Most organisations are pretty good at monitoring overall financial performance retrospectively on a periodic basis, for example through monthly management reports. They are not so good, however, when it comes to monitoring the factors that drive overall performance on a continuous basis^[1], which means that events happening within the business, whether related to problems or opportunities, are often not acted upon until it is too late.

There is also a growing need for information to be compiled and analysed that helps ensure that all elements of business, investment and other financial risks are being managed far more proactively. Recent events in the financial services sector that resulted in the economic downturn being experienced at the time of writing are perhaps the most graphic illustration of failure to track causes and effects adequately and can in such extreme cases lead to dire consequences. We come back to this later; meanwhile, let us note how such failures will almost undoubtedly result in yet more regulatory compliance requirements which will place ever more onerous demands on the information systems in place.

Looking on the positive side of cause and effect, however, another research study carried out by Freeform Dynamics^[2] identified an interesting linkage between effective use of business intelligence and organisational growth. In specific terms, a clear correlation was evident between those businesses that consider BI to be an enabler of competitive advantage and the likelihood that the business will be growing.

To turn this around again and focus back on the challenge, far fewer companies enjoy growth where business intelligence is not considered to be a competitive enabler, suggesting that BI is sometimes not afforded the priority it deserves (Figure 2).

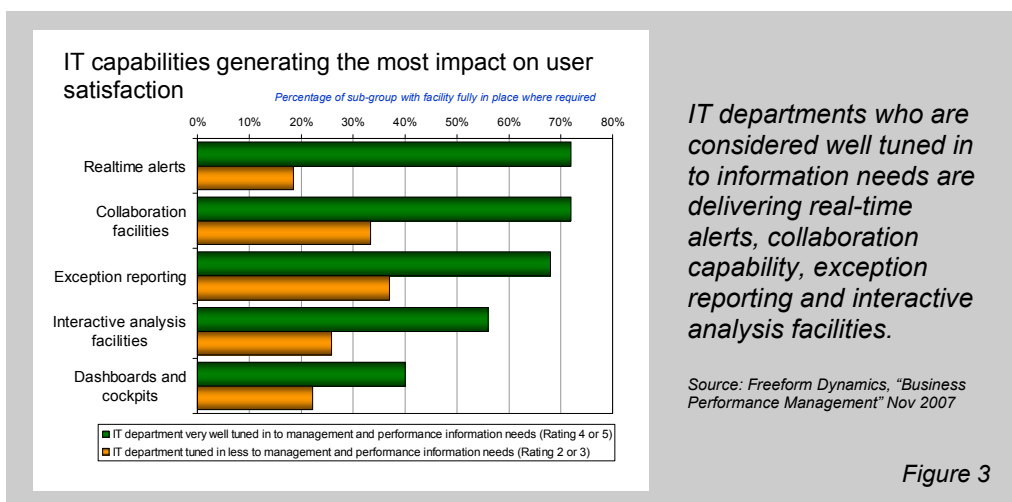


Just tell me what I need to know, when I need to know it

One of the pitfalls of delivering management information at a more granular level is that users can become swamped in data. Indeed, we saw in Figure 1 previously that 60% of respondents complained that the information they really require is buried in superfluous detail.

What these complaints translate into, is a requirement for intelligent and proactive delivery of information. Used correctly, business rules and other mechanisms embedded in data warehouses and other business intelligence components can draw on the wider information pool, to deliver a subset that is relevant to the user's needs, simultaneously drawing the user's attention to what matters the most and what is most likely to require action.

Taking pro-activity to the next level, selected information may actually be 'pushed' to the user in the form of status updates or alerts, delivered via email, a mobile device or other means: such mechanisms really do seem to be appreciated by users (Figure 3). Indeed today we are now beginning to see information being utilised directly in business processes in near real time to help line management make better, faster operational decisions.



And by the way, remember we are in an economic downturn

The last and probably most topical challenge given the backdrop of the economic downturn affecting organisations across the world is managing exposure and overheads. Indeed, it is clear that addressing the twin areas of 'risk management and control' along with 'containing and reducing costs' have become major factors in all areas of business decision making today.

Clearly, the fragmentation we have been discussing translates to both cost and risk. From a cost perspective, operational overheads in terms of systems management, maintenance, software licences

and power consumption associated with the proliferation of small-footprint BI solutions come into sharp focus. And when you consider that many of the systems in place, originally implemented to solve very specific requirements, are probably significantly under-utilised, the unnecessary overhead represented by an organisations 'BI solution landscape' could be significant. Add to this the productivity challenges arising from some of the issues we have highlighted and in the context of an uncertain financial climate, ways of tackling fragmentation start to look very attractive.

There is another way of thinking about the significance of BI in the current economic climate. With market conditions becoming both tough and unpredictable, the need for accurate information to steer the business on a more real time basis has never been greater. With this in mind, we can speculate that just as we have seen a direct correlation between effective BI and growth in the past, looking over the coming year or two, focussing on sound, accurate and timely intelligence could make the difference between commercial health and business exposure.

So, bearing in mind all of the above, what are the implications of the challenges and imperatives we have been discussing from a BI infrastructure evolution perspective?

Infrastructure evolution

Evidence from our research suggests that most reading this paper will recognise at least some of the above challenges and imperatives. Looking around at your infrastructure, however, probably utilising a mixture of IBM, Oracle and Microsoft database technologies running data warehouses and integration software on a variety of hardware, you may be wondering where to begin the journey to a more coherent and integrated business intelligence approach.

Possible approaches

At the highest level, assuming there is room for improvement and the will to act, there are three options:

1. Continue as you are, solving business intelligence problems at a local tactical level, in which case many of the problems identified previously are likely to persist;
2. Take measures to proactively rationalise and/or integrate existing data warehouses and other information stores into a logically and physically coherent architecture;
3. Implement a more gradual evolution programme, defining criteria and guidelines for new projects, effectively halting fragmentation and improving coherency over time.

The last option is really a "slow burn" version of the second, applicable when time or resources to spend on general modernisation and improvement are limited. For those who regard options 2 or 3 as desirable, and those that already have an improvement initiative in place, it is worth considering your situation and the options open to you.

We will look at some common scenarios in a little while, but first let us reflect on some of the key considerations and requirements for a future proof business intelligence infrastructure.

Key platform infrastructure considerations and requirements

When looking at such challenges, a clear set of considerations emerge with regard to the business intelligence infrastructure, particularly in terms of data warehousing platforms and associated integration and management tools. So let's now introduce this systems dimension into the equation.

Upping the level of currency and responsiveness

As we have seen, there is increasing evidence, both empirical and circumstantial, showing that the modern business must significantly change the way that it uses business intelligence / data warehousing solutions.

In the past it was commonly the case that the data warehouse was regarded as simply a 'big bucket' into which data was periodically placed, with some degree of data transformation and aggregation taking place along the way, usually via ETL (extraction, transformation and load) technology. This methodology has been in use for the last 10-15 years utilising reporting tools from organisations such as Business Objects (now SAP), Hyperion (now Oracle), Cognos (now IBM) and SAS (still staunchly independent). Such approaches have been the norm for management reporting. Indeed, many

enterprises have employed such static warehouses, typically combined with batch-oriented refresh and report generation, perhaps based on daily, weekly or potentially monthly cycles, as the foundation for supporting their business intelligence requirements. Whilst such systems are widely deployed, the value they deliver is variable and the growing demand for access to dynamic information means that for many organisations the static data warehouse approach may no longer be entirely suitable.

The evolving need to access business information on a more on-demand basis requires that data warehousing systems contain current, or at least near current, information. Equally it becomes necessary that reporting, query and analysis tools themselves can be used by decision makers in real time and that business processes be able to access such information in real time to enable greater business responsiveness. This move has placed new demands on the business intelligence infrastructure in terms of core warehousing capability to handle 'in flight' updates and refreshes of data, plus an increased emphasis on integration and BI process management automation that are essential to tapping into source systems in an effective and efficient manner. Even real-time 'drill down' from data in the warehouse into operational information stores has become a requirement.

In practical terms such requirements, coupled with the need to keep the data pool on which decisions are taken up to date and as free as possible from holding contradictory information, create a case for keeping warehouse systems in logically close proximity to as much of the source data as possible. In such situations it becomes possible to avoid the risks associated with the proliferation and movement of data, especially information integrity and security. In an ideal world one would want to access the information in place but this will not always be practical. However it is important that information doesn't undergo more transformation, aggregation etc. than is absolutely necessary. Multiple transformations of data can have unanticipated effects on data integrity (similar to the "Chinese whispers" effect), and it is also clear that in scenarios where the root data is held on disparate platforms it makes it more difficult to orchestrate and manage things centrally.

The requirement for information to be available to very large numbers of people and processes in near real time has in turn generated a number of technical demands on modern BI systems, including a clear need to process data in parallel from numerous sources coupled with the requirement for sophisticated, centralised metadata management and scheduling. These in turn require that adequate auditing, monitoring and staging capabilities, amongst others, be available and manageable from a central perspective.

Dynamic Data Considerations

Risks associated with data proliferation and movement to other platforms from the mainframe:

- *Integrity*
- *Security*
- *Access to Skills / Knowledge*
- *Cost of Management / Ownership*

Dealing with the access requirements of a broader audience

As already mentioned, there is a clear trend for formerly specialist Business Intelligence activity to align more closely with business job functions. One of the implications in infrastructure terms is for warehousing systems to be accessible from a range of familiar front-end tools, in addition to the more specialist analytical software that would typically be run by a business analyst.

Perhaps more important is the need for business applications and operational processes to utilise BI-generated information in real time. This necessitates creating systems that interoperate through a combination of service-oriented interfaces to information assets and open standards support, e.g. through published application programme interfaces (APIs), to allow easy and flexible integration with the user's tools of choice.

Data warehouse performance, scalability, resilience and security

As the volume of information routinely acquired by organisations grows and accelerates alongside the need to support many more users monitoring and managing business performance at a more granular level, there is a clear need to be able to cope with very large scale information access.

These factors combine to create the requirement that BI / data warehousing systems not only be capable of scaling to handle large data volumes, but also that they are highly performant, available on an 'ad hoc' basis and able to meet variations in demand in a predictable fashion. The watch-words here are performance, security, scalability and resilience. Gone are the days when the data warehouse being 'down' was just an inconvenience. With more organisations looking to embed BI-

enabled decision making into their day to day operational processes, the BI system itself becomes a business-critical component of the overall IT infrastructure.

Given the nature of the information being processed and the use to which such systems are put, the operation of these systems can become a matter of concern, especially as the requirements of their users, inside and outside the organisation, pose a wide range of security / data protection challenges. When coupled with the need to ensure data integrity and that access to sensitive information is restricted to those with 'a need to know', the operational security of BI / data warehousing systems becomes highly visible and sensitive. This consideration is especially pertinent when the use of the system is concerned with the management of risk in the business.

BI / Data Warehouse Evolution

Over the last few years database and business intelligence systems have slowly changed in their fundamental usage patterns and this is now driving architectural changes. Originally database systems created to handle OLTP (Online Transaction Processing) workloads were optimised to cater for real time access demands coupled with the need for them to be highly available, scalable and secure. On the other hand traditional data warehouse systems were optimised to handle read performance related issues and were typically addressed using complex data partitioning techniques coupled with sophisticated workload management tools.

These changing demands are reflected in the way that database / data warehousing / BI systems have developed. In the first phase of their use, 'Query' and 'Reporting' were the main functions delivered to provide, for example, periodic financial or sales reports. In the second generation of such systems the business demands were met via data mining type operations to perform multi-dimensional analysis to help optimise processes such as inventory management and merchandising. The latest business requirements focus around active business management and real-time decision making. This in turn promotes a need for a third generation of BI / data warehousing systems that can generate a wide variety of information, potentially from multiple data sources, when the user needs it.

The requirement today for systems that are capable of meeting the business needs highlighted in this document is forcing BI / DW architectural evolution. In essence systems capable of operating in a "dynamic data warehouse" environment must combine the characteristics from both the OLTP and traditional DW systems. However, it is difficult for systems built with read optimisation in mind to adjust to cater for the more interactive transaction processing type demands. Thus it is, at least architecturally, more straightforward for OLTP systems to handle sophisticated BI demands when built on a platform designed to handle these workloads.

This point quite nicely brings us on to the topic of the mainframe.

Mainframe Architectural options

The mainframe platform is still considered by many to be the "gold standard" server platform in terms of ease of management and, even more importantly has a reputation as the most secure computing platform available today, a fact enhanced by the robust management processes employed in its daily operations.

Ultimately, the common goal is to build a BI / data warehousing architecture that delivers the information business users require, when they need it, using a cost-efficient, manageable platform from which it is straightforward to address complex and sensitive security issues. The mainframe clearly holds great potential in such scenarios. Its security, scalability and operational efficiency profile make it a very attractive proposition when deploying dynamic data warehouse / BI systems, in the context of the discussion we have been having.

It is worthwhile noting that the mainframe has long been a cornerstone of business activity. Over its forty plus years of development the mainframe has become synonymous with availability and reliability. The fact that System z is designed for continuous operations is a characteristic that has, perhaps, fallen out of the minds of many. Unplanned downtime on System z is almost unknown and planned system interruption is in most organisations a very rare occurrence.

Such features become especially relevant when we consider the discussion around proliferation of small-footprint solutions, and the overheads these can cause. Considering all management aspects involved in BI / data warehousing, the low incremental Total Cost of Ownership (TCO) of the mainframe is likely to prove valuable.

Equally the very high levels of security protection that the mainframe routinely delivers will be valuable to many organisations, especially as data / information security is a matter of paramount importance to every business executive. The IBM System z9 and the IBM System z10 have long incorporated numerous capabilities to secure transactions and the information held on the platform. Capabilities include encryption, sophisticated identity management and authentication along with a host operating system that has been refined to handle the most complex and sensitive security requirements that extremely demanding customers can place upon it.

As we have already mentioned, this discussion is less about the 'absolute' value of mainframes versus distributed systems, and more about making the most of mainframe capabilities where they are present. As we explore these considerations further, we will be looking at three architectural options for implementing a forward-looking BI / DW strategy which takes into account the mainframe in general, and IBM's System z platform in particular.

The architectural development of the Mainframe as an Enterprise Platform

Before we examine the three potential architectural solutions utilised in BI / data warehousing environments it is worthwhile reviewing the architectural development of mainframe capabilities over the course of the last few years, several of which have potential use in such environments.

Many of the benefits enabled by the use of the speciality engines are described in more detail in the following section but it should be noted that in terms of operational licensing costs, mainframe offload engines such as zIIP, IFL and zAAP processors (see panel) effectively enable an organisation to build and operate an 'open' environment coupled to the mainframe within a single, secure and cost-effective System z platform.

It is also worth recognising that as workload demands increase the requirement for additional physical space, power supply and associated cooling scale much more slowly using a System z mainframe than alternative UNIX, Linux and Windows servers. The requirement for additional, expensive operational resources to manage the increased workload are also minimised as the mainframe was designed from its birth to be manageable as it scales. In fact published sources of data^[4] make a firm case that the mainframe is probably the most power efficient computing platform available today, which is an important consideration for organisations that have seen large cost increases for the electricity they consume, and even more so for those for whom getting adequate power into their data centres is a considerable challenge.

Offload Engines

Customers using System z Integrated Information Processors (zIIP) free up capacity in their System z9 general computing resources by effectively pushing eligible database oriented workloads (such as certain business intelligence and data warehousing queries) from the main System z9 processors onto dedicated zIIP engines. Whilst each zIIP engine needs to be purchased in addition to the System z9 processors, IBM has chosen not to impose any software charges on zIIP capacity. This licensing decision saves customers money compared with standard licensing models which impose software charges based on MIPS.

zIIP engines can also enhance business application performance. IBM DB2 for z/OS version 8 was the first IBM software fully capable of exploiting zIIP and workloads using select query processing (e.g. BI, ERP or CRM network-connected applications) and BI application query processing utilising the DB2 star-schema parallel query capabilities. These queries, along with DB2 utilities performing index maintenance structures, see significant performance and price-performance benefits.

The Integrated Facility for Linux (IFL) workload engine exists to run Linux applications on the mainframe. This option may prove to be attractive to organisations that utilize Linux applications as part of their BI systems but who wish to exploit the low cost of management and ownership associated with the mainframe along with the associated benefits of security and availability.

Lastly, we should introduce the System z Application Assist Processor (zAAP) processor, which handles Java workloads.

The Mainframe, Data warehousing and BI

Architecturally there can be considered to be three major elements to a BI / data warehousing environment:

1. Reporting and query interfaces;

2. Analytical engines and / or data integration engines;
3. Data store(s).

Experience has indicated that user interaction is ideally suited to graphical user interfaces. Thus, we can assume that the reporting and query interfaces will focus on the desktop or workstation.

The analytical engines provide a bridge between the reporting and query interfaces and the data stores. The choice of platform is dependent on two factors: firstly high speed access to the source data and secondly a cost-effective compute intensive environment. Today, all of the widely used solutions in this area easily connect with the underlying data stores, regardless of platform. Thus, the decision will typically be based solely on which platform provides the most cost effective operational profile for the organisation or on operational platform preferences / experience.

The selection of the optimal platform or platforms for data stores centres on factors that have already been introduced: speed, scale, security, reliability and cost. This decision can be more complex as several factors impact the choice of platform.

Taking the IBM System z as our reference point for mainframe computing, the architectural options considered as ways to meet future BI / data warehousing needs may be broadly classified for the purposes of discussion in this paper as:

- System z-centric BI / DW: locate the operational data store alongside the data warehouse on the mainframe;
- Hybrid BI / DW: mix data stores on the mainframe with those on distributed platforms;
- Non-System z BI / DW: store all data away from the mainframe.

We will now review the nature of each of these and the relevant merits of each model as the data store foundation for a forward-looking business intelligence infrastructure.

System z-centric BI / DW

In a System z-centric environment, all of the fundamental information components of the solution are hosted on the System z server, including the operational data store (ODS), the data warehouse itself and any data marts utilised in the system. This model is well-suited to address the strategic business needs described earlier in the paper centred around the demand for real-time or near real-time access to information to support key decision making or business process requirements.

Typically the systems will be based around the IBM DB2 offering functioning under the z/OS operating system. Any operational information sources not hosted on the System z platform are integrated using specialist connectors to provide a single hub of information utilised in the BI system. In this scenario the tools utilised by end users are either hosted on their workstations or run on distributed application servers.

The qualities inherent in Data Warehouse and BI systems utilising a 'pure' System z architecture are very much those intimately associated with the System z mainframe platform itself. Put simply the mainframe is usually regarded as the litmus test of servers when it comes to discussing factors such as security, scalability, reliability and availability. In addition, this model allows information to be accessed in rapidly varying business scenarios where key processes or decision making actions are based on having accurate and up to date information available.

Another factor to consider lies in the ability of BI / DW solutions based on System z to scale in a straightforward fashion. Several studies have shown^[3] that as System z solutions grow the management effort required to keep them operational does not scale in anything approaching a linear fashion. The manageability of System z solutions helps ensure that the cost of ownership of the mainframe compares very favourably with those of other platforms.

Studies have also shown that in terms of cost per user per year, System z can be lower to operate than other platforms^[4]. This factor has been greatly assisted over the last few years with the advent of mainframe speciality engines such as zIIP, IFL and zAAP (see previous panel). Given the increasing costs for the electricity to power IT systems and associated cooling, the cost of ownership of the mainframe has, if anything, become even stronger in the last twelve months.

It is worthwhile noting that IBM has also continued to invest heavily in the software capabilities necessary to integrate and cleanse data held in multiple systems as well as enhancing the ease with which reports can be built from the data pools. In many scenarios where the System z platform is relevant in the BI / Data Warehousing solution, it is important to ensure that the complexities surrounding data integration are addressed, especially in dynamic business situations. The mainframe has established great interconnectivity with a very wide range of alternative platforms and now boasts sophisticated software tools making it well able to address data integration questions. This is an area where the IBM Information Server Platform operates. On the reporting and analytics side of the equation the company has added to its existing capabilities through the acquisition of Cognos, whose software offerings now enhance the Information on Demand strategy that IBM has in place to help organisations better exploit real-time information.

A final factor is the ability of the System z to manage highly variable workloads in a secure fashion without imposing untenable demands on systems administrators. This matter is of no small consequence as data warehousing and BI systems are now forced to handle not only continuous, fairly predictable workloads generated by routine reporting but also on-demand analysis and management information generation on an ad-hoc basis by growing numbers of business users. The highly variable and unpredictable demand for service created by these mixed workloads plays well with the capabilities of System z, whose Workload Manager has been designed and refined over many years to ensure that access, control and service levels are maintained in line with defined business needs automatically as workload demands vary.

BI / DW on Hybrid System z Architectures

In so-called hybrid systems the operational data store and the data warehouse is hosted on the IBM System z Server running z/OS and IBM DB2. Statistical data marts and end user tools are operational on UNIX, Linux or Windows platforms. The end user systems may frequently only contain Web browser tools to provide access to the systems, possibly with local data stores and local manipulation / presentation products.

Hybrid systems share many of the benefits associated with the homogeneous architecture highlighted previously as much data storage and processing remains on the mainframe. The architecture comes into its own when the analysis tools that are utilised have not been written for the z/OS operating platform. In these scenarios the analysis applications continue to run on the platforms for which they were developed but can “transparently” make use of data resources held on the mainframe.

It is worth recalling how the mainframe platform makes use of speciality engines such as the IFL and the Internal Coupling Facility (ICF). The fact that IFL speciality engines allow Linux workloads to run on the System z platform without incurring any IBM software licence charges, potentially adds great attraction for organisations operating BI / data warehousing in these hybrid architectures. For example, the use of an IFL could permit an organisation to run any Linux dependent components of the BI solution on the mainframe with all of the attendant security, reliability, manageability, power consumption and associated TCO advantages therein, without losing the benefits associated with the use of such open source tools.

Why Choose a Hybrid Solution?

Hybrids involve users with data on distributed and mainframe platforms. Making the implementation mainframe-centric or distributed is a decision that organisations will base on a number of factors including:

- *Volume of data on the various platforms*
- *Skills associated with managing the data on the various platforms*
- *Security issues*
- *Capacity and scalability issues*
- *Availability of resources, power consumption, heat generation and cooling etc.*

BI / DW not utilising System z

In this architecture the System z platform plays no role at all, other than as a potential source of data to be accessed as required. These solutions typically utilise UNIX, Linux and Windows platforms and may also make use of dedicated data warehousing appliances. Such organisations may not employ mainframe systems at all in their environments.

Such systems can be straight forward to develop without the requirement to develop specialist mainframe skills. Until recently these ‘open systems’ solutions often benefited from speed of

development as specialist business analysis tools were developed on UNIX and Windows platforms. In addition, the capital acquisition costs associated with setting up open systems were often less costly than those associated with the mainframe and the skills necessary to operate open systems were readily available whilst mainframe administration skills may have been either unavailable or reserved for operating so called business critical applications.

It can also be argued that until recently, in particular with the development of the mainframe speciality engines, the mainframe platform did not offer similar levels of price performance to those achievable on open systems.

Before the advent of specialty engines it was difficult to cost justify a mainframe-centric approach to BI. Previously the visibly lower hardware and software costs of distributed environments, even when allowing for redundancy and higher management manpower costs associated with distributed server farms, attracted many organisations. It is possible that over time the advent of the IFL, zAAP and zIIP speciality engines may fundamentally revise the perceived price performance attractiveness of the mainframe when coupled with its security features and low TCO, especially as the mainframe continues to enjoy much lower operational costs, especially in terms of manpower and electricity usage, than other platforms.

The IFL Advantage

With the addition of the IFL the mainframe effectively becomes a logical hybrid architecture combining the best of Linux and z/OS in a single centrally managed system that may be a lower cost solution to traditional hybrid architectures.

Relevance of options

The important question to consider now is where and when the three architectural approaches are most appropriate given the attributes, benefits and limitations we have discussed. We will now look at each of these in turn then pull the various threads together into some recommended next steps.

Where is a mainframe centric approach to BI / Data Warehousing appropriate?

There are several factors that indicate when a mainframe centric approach to BI / DW should be considered. Situations that point towards a System z approach can be found below:

- Where significant sources of data (e.g. data warehouses, transactional, operational data stores etc.) are held in System z data sources including DB2, VSAM and IMS amongst others;
- There are existing System z / DB2 skills available and the organisation is prepared to continue to invest in them/expand them;
- Mission critical warehousing where 'Management', 'Security' and 'Risk' drive corporate policies;
- Organisations where System z is operationally connected to major data repositories;
- Scenarios with highly variable demand for standard reports, ad-hoc analysis, OLTP etc;
- Where continuous access to 'analytical data' BI / DW resources and reports is essential for people, other systems and business processes to operate effectively;
- Situations where the existence of multiple DW / Data Sources / Data Marts causes conflict.

Where is a System z / hybrid approach to BI / Data Warehousing appropriate?

- Keeping Data Warehouses updated in scenarios where the majority of source information is held on platforms other than System z;
- When geographic distribution significantly improves performance for users who are remote from the centralised mainframe resources;
- Scenarios that incorporate analytical tools that are not available on native z/OS;
- When a cost/benefit analysis determines that the complexity of a multi-platform environment is offset by the mixed price/performance profiles of the systems involved. In these situations it

is now possible that the use of mainframe IFL engines could provide an alternative to traditional hybrid approaches.

Where the mainframe approach to BI / Data Warehousing does not fit?

- Customers committed to utilising database platforms other than IBM DB2;
- Organisations where there is little / no System z experience / expertise;
- Organisations seeking to reduce mainframe workload.

Recommended next steps

There are essentially two main parameters to consider when looking at next steps – the current landscape in terms of existing data warehouse installations and the mix of sources being tapped into. These give rise to a range of scenarios, and recommended next steps for each (at least at a high level) are laid out in the table below. Where the scenario is against a green background, this indicates a good starting point for moving forward in strategic terms. An amber background means there may be some strategic adjustment required and red indicates a possible strategic misalignment that represents a potential future exposure of the business.

APPRAISING YOUR POSITION AND OPTIONS <i>The following is clearly no substitute for a full structured appraisal of your position and objectives, but identifying where your organisations sits in the table below may help in providing some high level pointers and ideas to ensure the optimum approach is being taken from the cost, value and risk perspectives.</i>			
LANDSCAPE / SOURCES	Operational sources mostly mainframe based	Mix of mainframe and distributed systems operational sources	Operational source mostly distributed systems based
<i>Mainframe currently the dominant architecture for data warehousing</i>	The architectural approach is almost certainly a good match for dealing with business requirements. Explore options for providing access via portals and desktop tools as necessary, and ensure you are up to speed with zIIP and the latest integration and management tools from IBM.	Make sure you are fully up to speed on the latest integration and management tools to streamline and effectively manage the integration of distributed as well as mainframe sources. There may be opportunities in both areas to boost service levels delivered to the business.	Make sure you are fully up to speed on the latest integration and management tool options to streamline and effectively manage the integration of distributed as well as mainframe sources. There may be opportunities in both areas to boost service levels delivered to the business.
<i>Equal dominance of mainframe and single distributed systems option</i>	There may be advantages in developing a strategy for a structured hybrid solution approach, with the mainframe acting as a coordinating hub. This would streamline integration with data sources and avoid disruption to existing local solutions,	The main imperative here is to tackle problems of fragmentation that are likely to exist in this kind of scenario. Again, there may be advantages to considering a structured hybrid approach with the mainframe acting as a coordinating hub to keep everything in sync.	The key to moving forward from this scenario is to evaluate how the mainframe and distributed warehouses are being used. It is common in this scenario for the mainframe to have assumed the role of a large legacy “bucket”, in which case a strategic review is required.
<i>Single dominant distributed systems option for warehousing</i>	This may be an acceptable situation, but the mainframe asset might not be fully exploited and addition of specialist offload engines could be useful. There may be advantages in exploring a strategy for a structured hybrid solution approach, with the mainframe as a hub.	The architectural approach is likely to be a good fit with business requirements. There may be advantages to considering leveraging the latest developments in mainframe warehousing and integration technology, but that depends on current capability and coherency.	The architectural approach is likely to be a good fit with business requirements. There may be advantages to considering leveraging the latest developments in mainframe warehousing and integration technology, but that depends on current capability and coherency.
<i>No real consistency in data warehousing usage</i>	This scenario represents a potential business risk for the future and a strategic review is highly recommended		

The role of Financing

Even before the economic downturn, IT projects needed to demonstrate either significant cost savings or show that they could deliver 'value' based returns on the investment being considered. Exploring ways of reducing internal IT costs is important, but significant work has already been undertaken to address this issue, although making greater use of the mainframe holds potential here too. When considering how more can be achieved to help business managers better exploit IT, making real-time operational information widely available offers new opportunities for the business itself to become more efficient and to find ways to increase the profitable revenue in tough market conditions. With this in mind it has been clear for some time that the availability of capital to invest in solutions, whether IT or anywhere in the business, has diminished. When coupled with the absolute requirement for any investment to be able to demonstrate a rapid return there has been a growth of interest in alternative sources of IT financing.

Existing sources of dedicated IT financing can provide traditional, offerings such as leasing but can also on occasion offer full blown project financing potentially covering all of the hardware, systems software, applications and services necessary to deliver a solution. Using IT financing can enable valuable projects to be undertaken even when macro economic conditions are tight and organisations, and their traditional banking partners, are reluctant to invest much capital into projects. The mainframe has long benefitted from there being a range of financing models available to help projects get off the ground, or indeed to cover their entire life spans. Even when mainframe servers are purchased directly, IBM has a tradition of ensuring that the engine investments are protected at future upgrade points, thereby helping the TCO case for the platforms use become even stronger.

Discussion

There are many organisations that today employ IBM System z mainframe computers to run many facets of their IT operations. The platform has merits that are widely, if often now sub-consciously, accepted including the ability to manage widely variable workloads with relative ease, the platform's outstanding security features and ease of security administration. Beyond this the mainframe's ability to scale without needing significant, additional, skilled management resource is perhaps less widely recognised, but is becoming an attractive feature. Finally the efficient power usage of the mainframe is almost certainly the least recognised of the platform's many characteristics.

Organisations that already utilise IBM System z may want to look again at considering the mainframe for business intelligence and data warehousing, areas that until recently may have been considered inappropriate on grounds of cost or performance. The ability of the mainframe to operate in such environments has been significantly enhanced over recent years as the IBM Information Server has matured. IBM has sought to make it possible to deliver, maintain and manage a range of BI / data warehousing environments employing a variety of server architectures, driven by its Dynamic Data Warehouse initiative which was designed to ensure that the mainframe had the functional capabilities required to handle parallel data processing, meta data management, scheduling, auditing etc.

Such features are, of course, essential to building a responsive, real-time BI system that can be administered without the need for additional operational resources. The addition of the IFL engine allows organisations to now build a logical hybrid BI solution entirely within the IBM System z9 and z10 platform combining the best characteristics of Linux for z and mainframe operations in a single, well managed and secure platform.

Mainframes can be cost effective: the market perception is that the mainframe is expensive to manage yet the reality is almost entirely the opposite. Once a critical workload level has been reached, additional workloads can be supported with minimal additional resources. This means that, while mainframe systems are not cheap in capital terms to acquire, the operational cost per workload unit drops consistently as the quantity of workloads increases.

It is clear that use of the mainframe is growing strongly in general terms; it is also evident that BI and data warehousing on the mainframe are important areas and that there is much to encourage organisations to employ such a platform as a major BI component. With IBM actively encouraging partners to get involved in creating and supporting such solutions, organisations would do well to review whether the mainframe can indeed play a part in their own BI strategies.

References

- [1] Business Performance Management, Freeform Dynamics, November 2007
- [2] The BI Inflexion Point, Freeform Dynamics, June 2007
- [3] IBM Scorpion Studies for Server Consolidation, 2000
- [4] Arcati Research Limited, The Dinosaur Myth, 2004

About Freeform Dynamics



Freeform Dynamics is a research and analysis firm. We track and report on the business impact of developments in the IT and communications sectors.

As part of this, we use an innovative research methodology to gather feedback directly from those involved in IT strategy, planning, procurement and implementation. Our output is therefore grounded in real-world practicality for use by mainstream business and IT professionals.

For further information or to subscribe to the Freeform Dynamics free research service, please visit www.freeformdynamics.com or contact us via info@freeformdynamics.com.

About IBM



At IBM, we strive to lead in the invention, development and manufacture of the industry's most advanced information technologies, including computer systems, software, storage systems and microelectronics.

We translate these advanced technologies into value for our customers through our professional solutions, services and consulting businesses worldwide.

For more information on IBM, please visit www.ibm.com.

Terms of Use

This report is Copyright 2009 Freeform Dynamics Ltd. It may be freely duplicated and distributed in its entirety on an individual one to one basis, either electronically or in hard copy form. It may not, however, be disassembled or modified in any way as part of the duplication process.

The contents of the front page of this report may be reproduced and published on any website as a management summary, so long as it is attributed to Freeform Dynamics Ltd and is accompanied by a link to the relevant request page on www.freeformdynamics.com. Hosting of the entire report for download and/or mass distribution of the report by any means is prohibited unless express permission is obtained from Freeform Dynamics Ltd.

This report is provided for your general information and use only. No warranty or guarantee is provided by Freeform Dynamics Ltd with respect to the suitability of the information provided within this document for any particular purpose.