IBM Information Management software

Exploiting DB2 on System z to meet your data warehousing needs

by Jay Yothers, Jaime F. Anaya, Terrie Jacopi, Oliver Draese, Wolfgang Hengstler and Anja Nicolussi

Contents

Executive summary

Enterprises are increasingly harnessing the power of business intelligence to gain a competitive edge. Today's business intelligence systems incorporate large data warehouses that are consolidated with near-real-time operational data stores (ODS) and continuously updated from multiple sources, such as financial, marketing and inventory databases. Often, thousands of users across an enterprise access a data warehouse with various business intelligence applications. These applications analyze and synthesize data into real-time insights that support fast and informed decisions.

It is therefore imperative that current business intelligence solutions cope with the emerging issues that arise from a mix of the following four workloads, including:

- *Continuous (near-real-time) data loading—similar to an OLTP workload.*
- *A large number of standard reports.*
- *An increasing number of true ad hoc query users.*
- *An increasing level of analytics and business intelligence-oriented functionality in OLTP applications.*

According to Gartner, these four workload types are increasingly top challenges for organizations beyond the growing size of the data warehouse, and mixed workload performance is well on its way to becoming the single most important differentiator issue in data warehousing. [1]

This white paper examines the role of data warehousing and DB2® on System z™ environments. It explains the benefits of running existing data marts and data warehouses on well-established System z data serving platforms to help facilitate real-time business intelligence analytics that meet today's requirements.

Why is data warehousing on System z the right choice?

In the past few years, IBM has focused heavily on achieving high levels of data warehousing sophistication and performance on System z. These efforts have received considerable notice across the industry. In fact, Gartner[1] ranks the System z as a data warehouse DBMS Server in its top leadership quadrant and states that a DB2-type solution is applicable to most mainstream data warehousing deployments.

There is substantial benefit when running data warehousing/business intelligence and OLTP on the same platform: DB2 for z/OS® not only runs enterprise OLTP systems, but also handles the data warehousing system. Thus, by standardizing and consolidating data marts and data warehouses on System z, organizations can leverage existing skills.

In addition, the workload management capabilities of DB2 for z/OS are tightly integrated with the Workload Manager for z/OS—thus providing the considerable strength that differentiates the DB2 for z/OS DBMS system. The Workload Manager provides optimum usage of the system resources, with the advantage of optimizing performance for mixed workloads, prioritizing jobs and query executions. Thus, for System z, the mix of Gartner's four distinct workloads is no problem at all, whereas on other systems, they will pose substantial challenges.

Other benefits include:

- *The ability to provide reliable, available, secure and protected access to data, which can greatly increase an enterprise's ability to fulfill compliance regulation.*
- *Hardware compression capabilities to help minimize disk space and augment I/O performance.*
- *A portfolio of data warehousing tools and applications on System z.*
- *The recent introduction of the zIIP specialty engine, exploited by DB2 for z/OS for complex query evaluations like star joins on complex database structures, to help free up capacity, optimize the System z IT infrastructure and drive down the price point of data warehousing workload on System z.*
- *The extension of DB2 for z/OS to support data warehousing with new functions like enhanced index options, new partitioning options and advanced SQL.*
- *Integration of DB2 for z/OS with the IBM Information Server Solution to provide high-volume, complex data transformation and data movement used for standalone ETL or as a real-time data transformation engine.*

The architecture behind data warehousing on System z

In July 2006, IBM expanded its data warehousing strategy to include the ability to deploy data warehousing applications on DB2 for z/OS and System z with the following:

- *DB2 for z/OS V8*
- *DB2 Alphablox*
- *IBM Rational® Data Architect*
- *IBM WebSphere® DataStage*
- *Tivoli® OMEGAMON® XE for DB2 Performance Expert on z/OS*
- *System z9™ Integrated Information Processor (zIIP)*
- *WebSphere Information Integrator Replication Edition*

While determining which business applications run on System z is a valid question, a more critical question is: Which applications run against the data on DB2 for z/OS?

Business applications run on their own application server environment. This application server is distinct from the database server where the data resides. The connection between these servers usually makes use of DRDA over a TCP/IP network. A particular business intelligence application may run on an application server that resides somewhere besides z/OS such as on Linux®, UNIX® or Microsoft® Windows®. As long as the SQL functionality is supported, the business intelligence application can execute on whatever application server meets the need of the application.

Pure System z business intelligence solution architecture and hybrid data warehousing on System z architecture

Digging deeper into variances of architecture for pure and hybrid System z business intelligence solutions, in a pure System z business intelligence solution as well as in a hybrid business intelligence solution, the System z environment is broken down into two parts: the operational data store and the analytical data warehouse. This means:

- *The **DataStage** loads operational data into the analytical.*
- *The **QualityStage** can also be used to cleanse and address any data quality issues, a function critical to the success of any data warehousing initiative.*

In addition, auxiliary data structures, such as materialized query tables (MQT) on System z further enhance query and analytics performance.

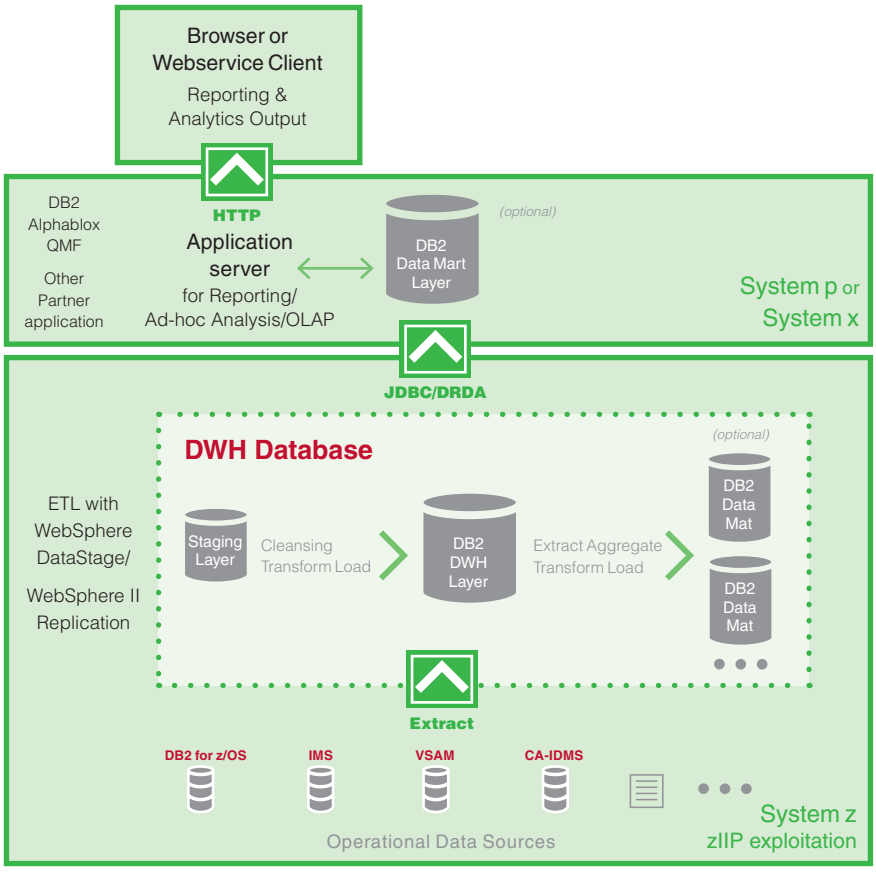

The hybrid BI solution approach

In a hybrid business intelligence solution, reporting is executed on distributed Web application servers, such as:

- *Alphablox*
- *QMF*
- *Cognos ReportNet*
- *Business Objects Server with relational, multidimensional (OLAP) and statistical data marts on distributed platform supporting end user tools, e.g. Hyperion Essbase, Cognos PowerPlay, SPSS.*

If reporting is done with DB2 Alphablox, the data is directly accessed on the System z through JDBC and queries take advantage of any available performance-enhancing MQTs. Business users can interactively access the Alphablox reports directly from their Web browser. But Alphablox also can be embedded as part of applications that are used in the normal course of a business user's job. This makes the analytics available "in line" with the business users' tasks. The information is available when and where it is needed to take action based on the extracted and aggregated business insights.

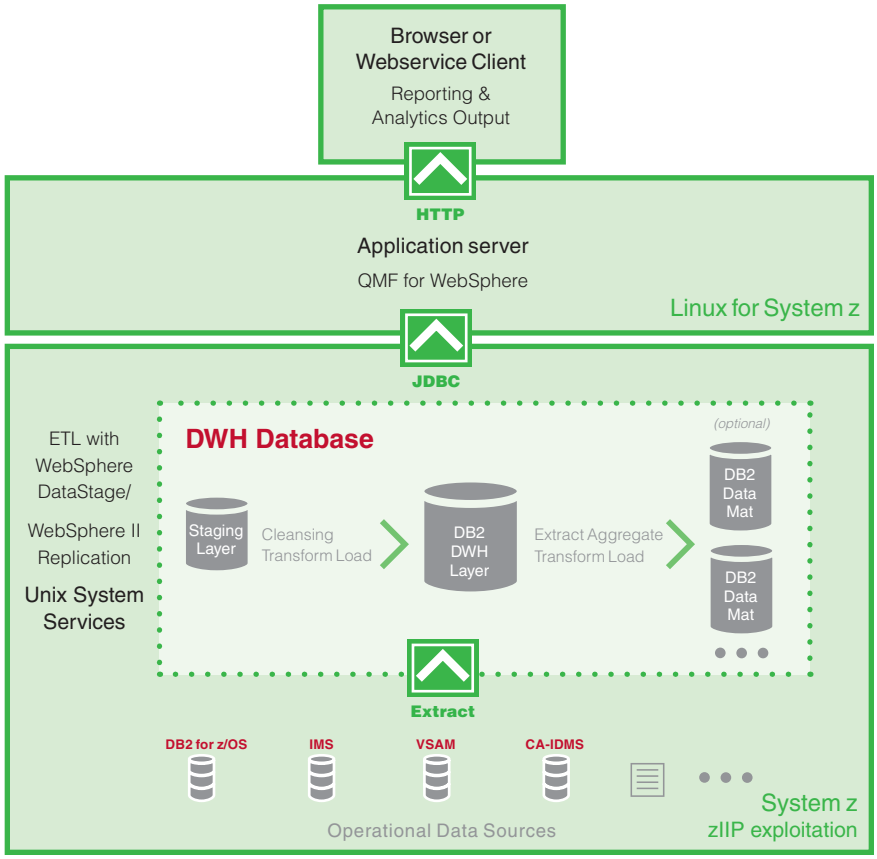

The pure System z BI solution approach

Why should data warehousing/business intelligence be done on the same platform as OLTP?

As data warehouses focus more and more on real-time orientation, SAP, Siebel and PeopleSoft – applications that already provide both OLTP and traditional data warehousing capabilities in a tightly integrated manner – are more inclined to run the application package on one platform.

At the same time, the market is shifting toward functionality in OLTP workloads. Because of time-to-market, competitive and improved customer insight, analytic results based on operational data or near-real-time data embedded in reports and ad-hoc queries are of high value.

The difference between OLTP and data warehousing/business intelligence is increasingly vanishing as more OLTP applications need advanced reporting and data warehousing/business intelligence require more recent data. By having all data on one platform, the constant need to refresh data marts with different processes is obsolete. The data mart content can be more recent than if it is extracted. Having different environments for the OLTP and the data warehousing/business intelligence data drives TCO up and allows more possible points of failure or intrusion.

An outcome of this market demand is a requirement for close proximity of data warehousing applications and online transaction applications. Enterprises that have an existing DB2 for OLTP in place and already use System z as their data serving platform are well positioned to consolidate business-relevant data marts on the System z platform.

A number of System z customers have already consolidated their data warehouses and data marts on System z, thus successfully exploiting DB2 on z/OS as a data warehousing/business intelligence platform.

That success is largely due to a solid foundation based on a tight integration that System z and z/OS share with the database (DB2 for z/OS) to provide unparalleled strengths for data serving on the mainframe. DB2 for z/OS exploits the System z platform strengths that have been there from day one. In addition, DB2 for z/OS continues to be enhanced, specifically for data warehousing and business intelligence functionality.

The top four emerging warehousing issues

As stated earlier, Gartner recently identified four workload types that represent the biggest challenges for organizations today. These include:

- *Continuous (near-real-time) data loading—similar to an OLTP workload.*
- *A large number of standard reports.*
- *An increasing number of true ad hoc query users.*
- *An increasing level of analytics and business intelligence-oriented functionality in OLTP applications.*

The next few sections take a closer look at these four workload types.

Continuous (near-real-time) data loading

Traditionally, data warehouses extract data from the operational storage through batch processes executed within maintenance windows. In modern data warehouses, however, the need to obtain the most recent data is becoming more important. The data needs to be extracted and aggregated more often to be sufficiently recent, and the complete process needs to handle bigger amounts of data in less time. The process of extracting data from the operational storage, transforming it to the necessary aggregates, moving it to the target database and loading it again, is known as the ETML process. This process also needs to be flexible enough to add different data from the operational storage to the data warehouse. This might be necessary because the granularity of information in the warehouse is not well suited for ad hoc queries or because new reports simply require new kinds of information.

Performing the ETML also needs to be done in as little time as possible to allow these queries to operate on the most recent data. It is therefore preferable to optimize the process and to avoid unnecessary steps.

The process of extracting information from the OLTP system by necessity runs on the same platform as the OLTP system. The process of transforming this data to conform to the needs of the data warehouse is also usually performed on the platform where the OLTP system resides. This is because many of the techniques used in the transformation process reduce the amount of data that then must be moved to and loaded into the data warehouse.

Regardless of where the transformation process is performed, it is necessary to move the data to the platform housing the data warehouse. There are many methods that can be employed in this movement of data. Whatever method is employed, if the data warehouse is physically distinct from the platform housing the OLTP system, it is necessary to transmit this data over a communication path of some kind.

This can be a time-consuming operation that can be removed completely if the data warehouse resides on the same platform as the OLTP system. With the Parallel Sysplex® capabilities of the System z, it is a relatively simple matter to have the data warehouse on the same platform as the OLTP system and, at the same time, be on distinct processors or machines from the OLTP system.

If extraction and transformation takes place on the same platform as the OLTP workload, it is important to assign only defined resources to this process. The appropriate assigning of defined resources prevents the transformation of the data from taking up too many resources and thus slowing down the main (OLTP) workload. This assignment of resources to OLTP and other resources to the ETML process can be achieved with the help of the Workload Manager (WLM). The Workload Manager component of z/OS has proven its worth in its ability to maximize use of available resources. WLM manages widely varying workloads efficiently and effectively, allowing you to fully utilize the systems resources you have available. This means you can run your data warehouse workload together with the transactional (OLTP) workload on the same DB2 subsystem or different DB2 subsystems on the same system.

With WLM, a particular element of work can be given an initial priority based on business needs. More importantly, over time, the priority of a given element of work can be altered based on changes in business needs as expressed in the WLM policy. DB2 for z/OS works hand-in-hand with WLM to ensure that these priority alterations take effect immediately with respect to query processing, regardless of how that query has been parallelized. Capacity can be brought to bear whenever it becomes available.

Large numbers of standard reports

The number of reports in a data warehousing or business intelligence environment is continuously growing. Instead of having only special reports for analytical purposes, the amount of standard and "every day" reports are increasing dramatically. Customers require analytical information on demand; business performance metrics have to be presented in an ongoing manner rather than just periodically.

Not only do reports need to be created, run faster and process more data, but the data itself needs to be at a more detailed level. Even though reports may have a common base, most reports have the need for "special" detailed data to process. So this data must either be moved to the data warehouse entirely or queried as needed from the operational storage.

The second option allows users to work with more recent data while avoiding the risk of having several copies of the same information, which might get out of sync. In addition, a mature workload management system is important to prioritize the normal OLTP workloads over the report generation to guarantee that operational tasks will be performed correctly.

An increasing number of true ad hoc query users

With the growing complexity of available data, it has become difficult to predict which information is or will become relevant to analytical users. It has become nearly impossible to define what needs to be transferred into offloaded data warehouses in advance. So to enable ad hoc queries that require real-time information, all data needs to be available to the analytical users all the time without the need to request and to wait for it. The data needs to be accessible 24x7 in much the same way as for normal OLTP usage. And security needs to be enforced so that everyone with a "need to know" can access specific data and query the appropriate information without other data being compromised.

The IBM System z hardware, z/OS operating system and DB2 for z/OS are designed with reliability characteristics like self monitoring, redundancy, self healing and dynamic configuration and management. For example, in DB2 you can make database changes, such as adding a partition, without an outage.

With IBM's Parallel Sysplex and DB2 data sharing solution, the System z platform provides the highest level of availability possible. It provides the ability to do rolling upgrades of service or versions of the software stack without any application outage. It also provides users with the ability to grow their processing needs horizontally by adding systems to the Sysplex to accommodate their growing OLTP or business intelligence needs – again, without taking the system down or requiring application changes to ensure performance.

Expansions to security allow for row-level granularity, helping with the security issues of Web-related applications as well as business intelligence and data warehousing applications. For additional protection, users can encrypt their warehouse data while it sits in their DB2 system.

An increasing level of analytics and business intelligence-oriented functionality in OLTP

In many applications, the user is not only presented with a single data record, but also with summaries and analytical output. This mixture can be handled by parallel access to multiple data layers where some information is extracted from the OLTP store and others from data warehouses. But the warehouse data is normally not as recent and accurate as the OLTP data. So even contradictory data might be visible to the user within the same application where the older warehouse data doesn't represent a correct summary of the shown OLTP data anymore.

The only solution is to keep the warehouse data as close as possible to the original OLTP data or even to perform all queries on the operational data store directly.

The usage of DB2 materialized query tables within the ODS delivers the data warehouse capability with pre-processed OLTP data as summary. The result of a complex summary or analytical query is not disposed after the query results have been fetched, but it is stored within a table again. The same queries can be executed repeatedly just as on data warehouse data without the need to recalculate the summary information. The DB2 optimizer even takes these MQTs into consideration to create faster response of different queries if the MQT data can be partially reused to resolve the result set. At the same time DB2 ensures that the materialized data remains actual enough to be used by the optimizer. If the data exceeds a specified age, it is ignored by the optimizer. The administrator then simply has to trigger a refresh of the MQT content.

Regulatory compliance

Several regulations like the Sarbanes-Oxley Act (SOX), Basel II, Data Protection Act (UK) or the USA Patriot Act were created to protect investor's interests, reduce fraud and improve reporting of financial data.

Companies have to be able to prove that their data is safe, that all changes to data are registered and that even in case of a disaster, no information, which is necessary to run the business, is lost. This implies that a DBA has to ensure:

- *A proper backup and restore concept.*
- *That access to data is only granted to the persons with a need to know.*
- *That all changes are logged and can be reviewed for several years after the change.*
- *That there is no single point of failure.*

In addition these tasks, the regulations also require that a high quality of data is reproducible, so that there aren't several versions of the truth, caused by slightly different versions of the data.

The System z platform meets the highest industry security certifications. Features to support data encryption are built-in, even at the hardware level. Authorization is implemented as low-level operating system functionality, delivering the ability to select which user can access which portion of data. Detection services are designed and built in to prevent unauthorized access to data and to report all intrusions. The hardware-supported encryption is used to ensure that backups are not readable if they fall into the wrong hands. With DB2 9 for z/OS, the network communication between database and client can be SSL encrypted based on network-trusted context. Numerous IBM products are designed to support regulation goals, such as:

- ***DB2 Audit Management Expert****, which helps administrators collect all information necessary for an audit. It maps the audit requirements and information requests to the underlying technical structures and delivers the information without having the need to collect all the parts from the several different points within the system manually.*
- ***DB2 Data Archive Expert****, which supports administrators meeting compliance regulation requests by moving rarely used data to cheaper storage systems to ease the process of retaining all data for longer periods of time.*
- ***IBM Encryption Facility for z/OS****, which helps ensure that information is stored encrypted on all types of storage media, so that even if a physical intrusion occurs, all of the gathered information is still protected. This also allows the safe sharing of information with partners or customers.*
- ***DB2 Test Database Generator****, which helps application developers needing sample databases to develop their applications but who don't have the permission to view actual data from the operational database. This tool generates test data from scratch or from existing data of the production environment by removing all sensitive information.*

Finally, by keeping the warehouse data on the same platform as the operational data, the administrative overhead for audit management can be kept at an optimal minimum by avoiding several data sources which need to be reconciled or reported separately. There is only one central data repository which needs to be compliant instead of multiple systems which each need to be handled separately. The above mentioned tools ensure that audits on the System z platform can be done with minimum cost and optimum efficiency.

Technical considerations

The following section lists several technical aspects that show why System z, z/OS and DB2 for z/OS are outstanding choices for data warehouse management. DB2 for z/OS has been enhanced significantly in versions 8 and 9 with regard to data warehousing techniques, allowing query execution and data management to be done faster and with less TCO. At the same time, the ability to access and handle huge amounts of data has been dramatically improved by, for example, increasing the amount of possible partitions from 254 to 4096.

The z Integrated Information Processor (zIIP)

Designed to maximize resource optimization, the zIIP processor is priced less than general purpose processors, and the MIPS it provides do not count toward the software costs of the system. DB2 exploits zIIPs to execute most of the DRDA tasks from distributed workloads. Communication with DB2 over the network, e.g., to move large amounts of data, doesn't need to be handled by the general purpose CPs anymore but can directly be executed on zIIPs.
To execute queries in parallel is a very common use in data warehouse environments. The coordination of parallel execution utilizes the zIIP too, taking workload off the main processors and reducing the costs.

The DB2 utilities also exploit this specialty processor by having major operations like sorts being performed on the zIIP engine, sparing the general purpose CPs.

Parallelizing of queries

DB2 exploits different types of parallelism, including:

- ***Query I/O parallelism****, which manages concurrent I/O requests for a single query, fetching pages into the buffer pool in parallel. This processing can significantly improve the performance of I/O-bound queries.*
- ***Query CP parallelism****, which enables true multitasking within a query. A large query can be broken into multiple smaller queries. These smaller queries run simultaneously on multiple processors accessing data in parallel reducing the elapsed time for each query. Starting with DB2 V8, the parallel queries exploit zIIPs when they are available on the system thus reducing the costs.*
- *To further expand the processing capacity available for processor-intensive queries, DB2 can split a large query across different DB2 members in a data sharing group, known as* ***Sysplex query parallelism****.*

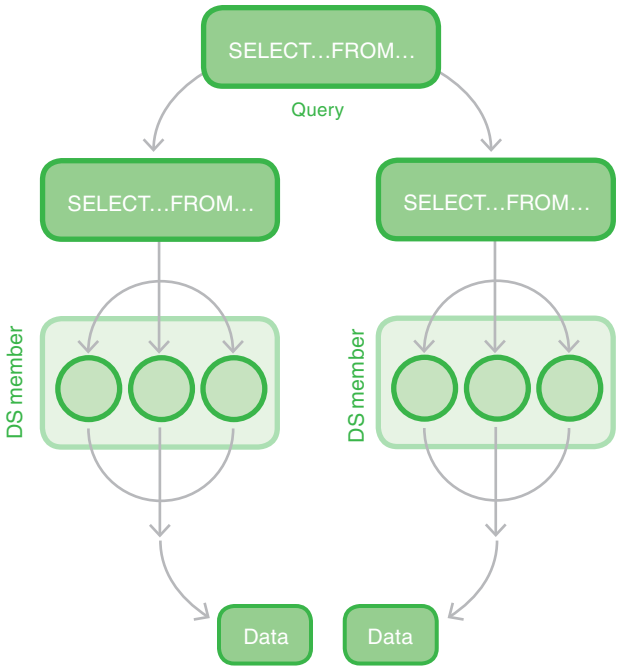

Each query is different and requires different degrees of parallelism. DB2 can determine the optimal degree of parallelism to use for a given query and utilize CPU parallelism within and across the optimal number of members of a data sharing group to achieve the best response of a given query.

DB2 for z/OS also supports parallel utilities. This means data can be loaded into different partitions in parallel. DB2 provides a simple range partitioning scheme to make this a trivial exercise.

Hardware compression

Hardware compression is a technique where data is compressed by the database so that it occupies less disk space. Compression and decompression of data are accomplished on the System z environment via a hardware instruction, making it faster than the software-based compression algorithms used on other platforms. As each generation of System z processor gets faster, the compression feature will also benefit from this development.

Compressing data reduces elapsed time of most data warehouse type queries. DB2 for z/OS compresses all data of a complete page. It uses the hardware instruction along with a data dictionary to give the most efficient compression available.

With a rule-of-thumb 50% compression rate, a compressed page contains twice the rows an uncompressed page would contain. This means that each I/O retrieves twice as much compressed data as it would if the data was uncompressed. The data remains compressed in the buffer pool, which means that DB2 for z/OS can cache twice as much data in its buffer pool. Finally, when data is modified in a row that is compressed, the information logged about that data change is also compressed, thus reducing log volume.

The larger amount of data retrieved in each I/O gets compounded with the DB2 for z/OS pre-fetch mechanisms, which provides significant elapsed time reductions for all sequential processes. This includes the typical business intelligence queries that make use of table scans and index range scans. However, this also includes other sequential processes such as utility access, providing additional benefits in terms of faster reorganizations, faster unloads, and faster recovery.

Star join enhancements

In many business intelligence environments, the data is split into multiple dimension and few fact tables. Performing queries on this information makes the join over many tables necessary where dimension tables may have significantly fewer rows than the fact tables. Starting with DB2 V8, several enhancements have been made to accelerate the execution of star join queries.

Beyond increasing the limit up to 225 tables within joins, performance was optimized by using "In-Memory Workfiles" (IMWF), which contain the join predicates and selected columns only, enabling binary search. The optimizer has also been enhanced and can now predict the costs for joins over dimension tables in a more precise way.

Materialized Query Table (MQT)

Expensive calculations need to be performed to add information to the base data, to create summaries or aggregates or even to join the data of several tables into a bigger, denormalized form again. MQTs can help avoid the repeated execution of these tasks. Once processed, the information is stored in a table again, accessible for later use. This way, lots of processing power can be spared by reusing previous query results.

Not only do direct queries against the MQTs receive the benefit of these materialized results. The DB2 optimizer uses the MQTs (partial or as a whole) to accelerate different queries. Otherwise these queries would have to perform subtasks of the same processing to create their result. This can be done because MQTs are more than just stored query results. The operation and its access path are maintained as well, which allows an easy refresh of the MQT content without having to specify the source query again.

Summary

In light of the arguments presented above, we highly recommend the consolidation of infrastructure based on DB2 for z/OS and System z on an already existing System z-based data serving platform.

For more information

To learn more about leveraging the unique System z differentiators to meet your business intelligence and data warehousing needs, please visit **ibm.com**/systems/z.

Additional resources

As part of its efforts to support ongoing research and development, IBM has established a best practices group. This team consists of experts from DB2 on z/OS development, professionals from technical sales, field sales teams with long-standing customer experiences in data warehousing and DB2 for z/OS. This group incorporates advice from marketing as well as from customers who have already trusted their business-critical information processes, related data warehouses and data marts on DB2 for z/OS.

Results of the group will be published regularly through User Group conferences such as the Guide, Share and System z technical conferences as well the DB2-related technical conference. In 2007, the team intends to publish "best practices" guidelines that will include recommendations regarding sizing, zIIP usage, workload, hardware and software configurations.

More information is available on the following Web sites:

- ***Announcement letter***
  ibm.com*/common/ssi/fcgi-bin/ssialias?subtype=ca&infotype=an&appname=iSource&supplier=897&letternum=ENUS207-042*
- ***System z***
  ibm.com*/software/data/db2bi/systemz.html*
- ***Why business intelligence on z9/zSeries Makes "Cents"*** *– replay*
  ibm.com*/software/os/zseries/webcast/26jan*
- ***Data warehousing on z..is it right for you?*** *– Teleconference in replay*
  ibm.com*/software/os/zseries/telcon/7sep/*
- ***IBM data warehousing and analytics***
  ibm.com*/software/data/db2bi/tab_bcu.html*
- ***Information Center for z/OS****: Information Management Software Information Centre for z/OS Solutions, information about DB2 for z/OS, DB2 QMF, IMS, and DB2 and IMS Tools from a single location*
  *publib.boulder.****ibm.com****/infocenter/dzichelp/v2r2/index.jsp*

IBM Corporation
New Orchard Rd.
Armonk, NY 10504
U.S.A.

Produced in the United States of America
04-07

[1]Gartner, Magic Quadrant for data warehouse DBMS Servers, 2006, Publication Date: 25 August 2006, ID Number: G00141428

TAKE BACK CONTROL WITH **Information Management**

IMW11861-USEN-01