Using IBM System z As the Foundation for Your Information Management Architecture

Redguides
for Business Leaders

Alex Louwe Kooijmans
Willie Favero
Fabricio Pimentel

- Information management challenges
- Strengths of IBM System z for information management
- Exploring the IBM Smart Analytics Optimizer for DB2 for z/OS V1.1

Redbooks

Executive overview

Many companies have built data warehouses (DWs) and have embraced business intelligence (BI) and analytics solutions. Even as companies have accumulated huge amounts of data, however, it remains difficult to provide trusted information at the right time and in the right place. The amount of data collected and available throughout the enterprise continues to grow even as the complexity and urgency of receiving meaningful information continues to increase.

Producing meaningful and trusted information when it is needed can only be achieved by having an adequate information architecture in place and a powerful underlying infrastructure. The amounts of data to mine, cleanse, and integrate are becoming so large that increasingly the infrastructure is becoming the bottleneck. This results in low refresh rates of the data in the data warehouse and in not having the information available in time where it is needed.

And even before information can become available in a BI dashboard or a report, many preceding steps must take place: the collection of raw data; integration of data from multiple data stores, business units or geographies; transformation of data from one format to another; cubing data into data cubes; and finally, loading changes to data in the data warehouse. Combining the complexity of the information requirements, the growing amounts of data, and multiple layers of the information architecture requires an extremely powerful infrastructure.

This IBM® Redguide™ publication explains how you can use IBM System z® as the foundation of your information management architecture. The System z value proposition for information management is fueled by the traditional strengths of the IBM mainframe, the specific strengths of DB2® for z/OS®, and the broad functionality of the IBM information management software portfolio. For decades, System z has proven its ability to manage vast amounts of mission-critical data for many companies throughout the world; your data is safe on System z.

The available information management functionality on System z has grown from database management systems to a full stack of solutions including solutions for content management, master data management, information integration, data warehousing, and business intelligence and analytics. The availability of Linux® on System z provides an excellent opportunity to place certain components in an easy-to-manage and scalable virtualized Linux server, while benefitting from the System z hardware strengths. DB2 on z/OS can remain the operational data store and the underlying database for the data warehouse.

The next generation of System z is growing into a heterogeneous architecture with which you can take advantage of System z-managed “accelerators” running on IBM System x® or IBM Power Blades. The first of these accelerators is the IBM Smart Analytics Optimizer for DB2 for z/OS V1.1, an “all-in-one” solution in which System z, z/OS, DB2 on z/OS, an IBM BladeCenter®, and IBM storage work together to accelerate certain queries by one to two orders of magnitude.

With the IBM Smart Analytics Optimizer, slices of data are periodically offloaded from DB2 on z/OS to the BladeCenter. After a query is launched against that data, it will automatically run against the data kept on the BladeCenter. The BladeCenter will process the query an order of magnitude faster than DB2 on z/OS, because all data is cached in internal memory on the BladeCenter and special compression techniques are used to keep the data footprint small and efficient.

As a solid information management architecture ready for the future, System z has it all.

- It has proven ability to manage massive mission-critical databases.
- It has superior availability and scalability characteristics, providing true data sharing technology.
- Virtual Linux servers provide flexibility and scalability.
- The large capacity of the System z servers, combined with the energy-efficient design, provide the ability to run a low-cost centralized information management infrastructure.
- A full-function software solution stack covering all layers of an information architecture
- In many cases, System z is already the data and information hub of the enterprise.
- Many transactional workloads already run on System z, creating and updating large portions of the enterprise’s operational data.
- It offers a future roadmap for a heterogeneous architecture, providing even more flexibility in terms of deploying workloads on System z engines, IBM System x blades or IBM Power Blades.
- A strong drive and vision by IBM to continue to deliver System z as the leading platform for enterprise-level information management needs.

Information management challenges

Efficient and effective information management is key to the success and competitiveness of an enterprise. Having the right information available at the right moment for the right people is critical for any business process. Simply designing and implementing smart application programs is not enough to achieve this goal. Instead, solutions are needed that organize information in such a way that it can be accessed when and where it is needed.

An additional upcoming issue, as discussed later in this document, is that IT infrastructure and IT resources can inhibit you from achieving this goal. This can sometimes occur because accessing information is so resource-intensive that even on the largest systems information cannot be returned in a timely manner. Also, accessing the same data from many programs at the same time can cause serious performance problems too, especially when updating data.

Having a powerful infrastructure is critical to implementing an effective information management architecture.

The information challenge

Today, information is one of the most valuable assets that a company possesses. Information consists of data, and data is everywhere, in large quantities. Unfortunately, the reality in many organizations is that data is spread across the organization, duplicated, conflicting, in silos, hidden, buried in an application, inaccessible, incomprehensible, compromised, unsecured, inaccurate, meaningless, out of context, and out of place.

Data is meant to provide insight into business, supplying a source for performance metrics and offering a competitive edge. But without a well-designed structure and a well-planned architecture, data becomes a huge challenge to contain, manage, and understand. It is almost impossible to derive useful information from meaningless data. But you will probably not obtain useful information from useful data either, without implementing a well-designed information architecture.

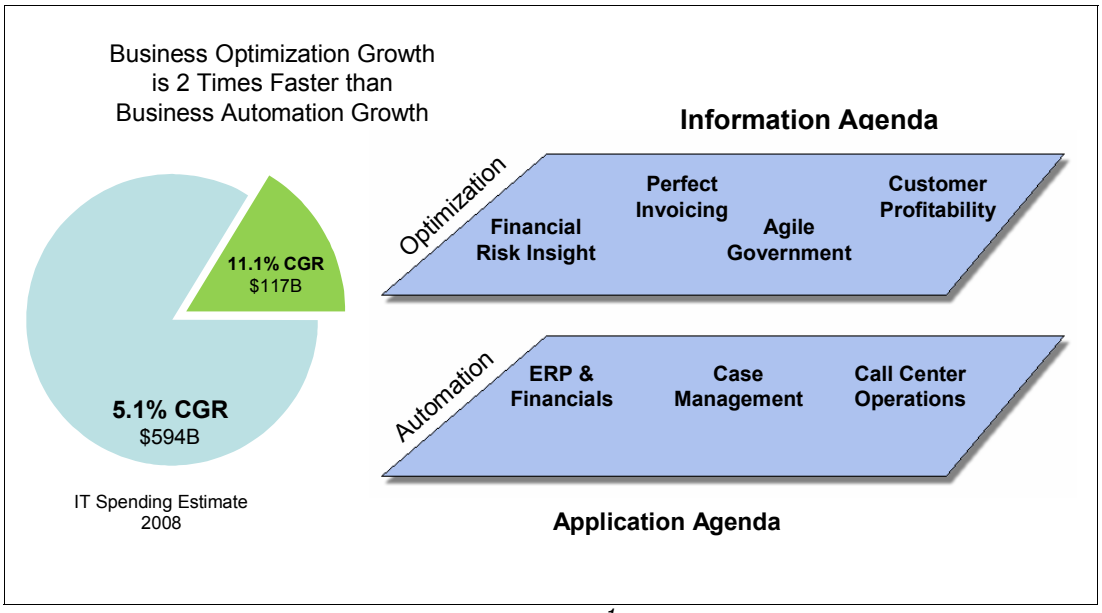

Figure 1 Focus on business optimization is accelerating[1]

Companies are investing more in optimizing business than in simplifying automation processes (Figure 1). They are investing to better use information, including applying better analytics to more information to drive better decisions and smarter business outcomes, and this is true across all industries (Figure 2).

[1] Includes hardware, software, and services. Does not include networking, printer, or stand-alone printer or PC markets.

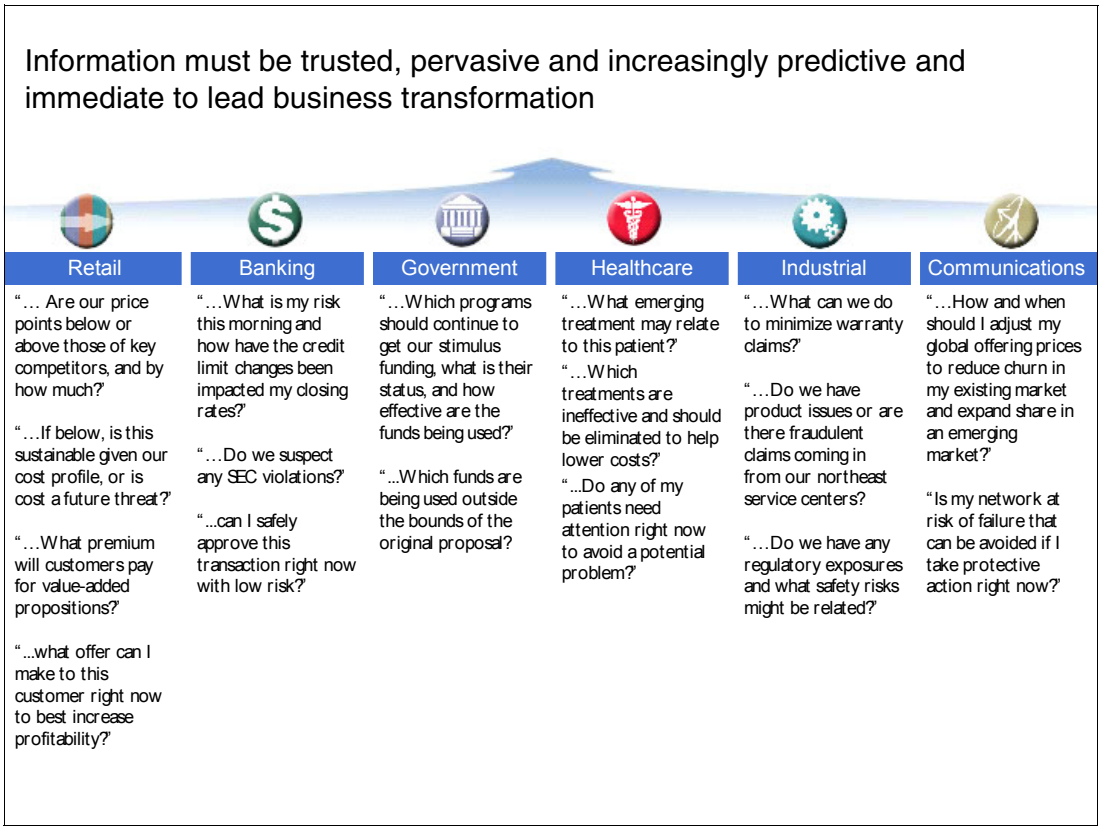

Figure 2 Industry needs for information

The idea is to create trusted information and ensure that it is delivered when, where, and to whom it is needed, thereby leading business transformations to better serve clients, patients, and citizens.

The ultimate challenge: coping with large quantities of data and information

Those who have finally reached the goal of structuring information in an acceptable manner are faced with the next challenge, which is managing large quantities of data and information.

Focus is shifting from building a functional information management architecture to providing an infrastructure that is powerful enough to run it.

Business intelligence and business analytics rely on the availability of massive amounts of data. This trend is continuing based on the growing importance of the following aspects:

- The need for reporting is increasing.

  Reports need to provide useful information, and they are increasingly based on the complex integration and mining of multiple data sources across the enterprise. The increasing need for reporting is a result of increasing needs from company executives to continuously be able to assess the state of the business. But there are also growing reporting needs as a result of regulatory requirements.

- The amount of data is simply increasing year over year.

  Enterprises tend to collect more and more data; consider of all the data that is collected when people visit web sites, let alone do business on the Internet. The way companies do business these days, especially using the Internet channel, warrants a steady growth of data, and the end of this growth is nowhere in sight.

- Being able to predict the future is becoming a significant issue for most companies.

  Predictive analytics is the capability to make decisions based on future trends. These future trends are derived from current trends and future parameters, within and outside of the company. To be able to detect current trends, massive amounts of data are again needed, and this data needs to be organized and analyzed. On a computer system, this is an I/O- and CPU-intensive process.

These developments require an intensive focus by IT executives and architects on both the information architecture and the underlying infrastructure. Using a few Windows® servers with several BI tools under the desk is an ineffective approach. In fact, the situation is much more serious; the growth of data, combined with the complexity of the business intelligence and analytics requirements, will become a major issue for many companies. Even now, in most companies data in data warehouses is at least 24 hours old, and in many cases as old as 7 days. Sometimes this occurs because there is no need yet for almost current data, but often it is because the infrastructure is not powerful enough to integrate, move, cleanse, and refresh data in a timely manner.

A study by Gartner has revealed the following information (Figure 3):

- More than 90% of Global 2000 companies plan to incorporate analytics into multiple operational applications that access the data warehouse by 2010, but fewer than 15% of data warehouses have been designed to provide high availability, failover, disaster recovery and the remaining components of mission-critical systems.
- By the end of 2009, 90% of Global 2000 companies will have implemented some type of mission-critical dependency between the warehouse and at least one revenue supporting or cost-controlling operational application. This is up from less than 25% in 2007.
- Fewer than 15% of data warehouses in 2007 were designed to provide high availability, failover, disaster recovery and the remaining components of mission-critical systems.
- The majority of companies using data from a data warehouse for inline, operational analytics have reported that a data warehouse production failure caused operational systems to cease daily operations, resulting in lost revenue or increased costs.

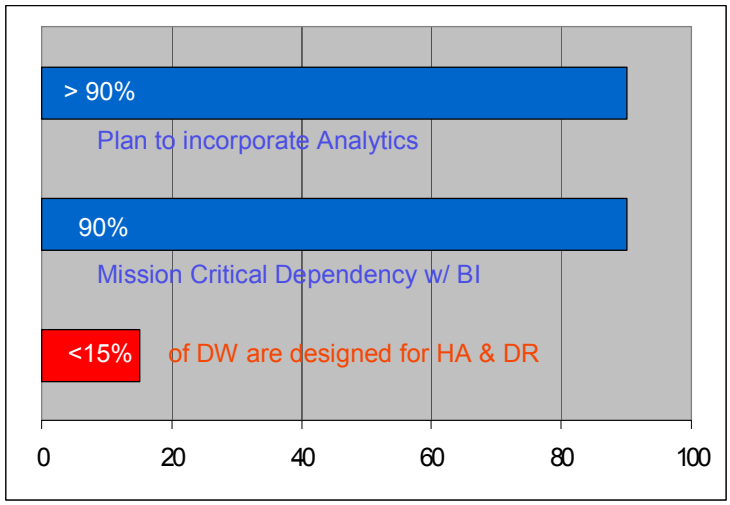

Figure 3 Operational analytics and the emerging mission-critical data warehouse, 14 May 2007

The importance and need for an infrastructure with the following characteristics is increasing rapidly:

- The ability to run databases with many terabytes of data in a trusted, secure and high-performing manner, with many programs and users accessing these databases at the same time.
- The ability to refresh data in the data warehouses at a refresh rate required by the business. The trend is that this refresh rate is increasing for certain business processes, so that data in the data warehouse becomes more current.
- The ability to perform complex data integration between multiple data stores spread across the enterprise.
- The ability to perform smart analytics on potentially all the data inside the company. In most cases these analytics run against the data in the data warehouse, but in various business processes the need for “real-time” analytics is increasing. Real-time analytics means that the analytics process runs at the moment a transaction takes place, for example a credit card transaction, and uses current data.
- The ability to deliver and infrastructure with all of these characteristics in a highly available, or even continuously available environment. Depending on the way availability is implemented, continuous and massive data replication can be required.

  Currently, by far most business intelligence and business analytics are batch-type activities that are not significantly impacted the system is not available for a few minutes or hours. However, with the increasing need for real-time analytics, the availability requirements are becoming more stringent, as well.

Later in this document we explain how IBM System z addresses these requirements with an additional boost using the IBM Smart Analytics Optimizer for System z.

The IBM Information Agenda

Information On Demand is a well-established term that has been used in information management for a few years now. The following four entry points have helped companies to start addressing their most pressing needs, while enabling them to embrace their long-term information management objective:

- Manage data over its life cycle
- Optimize content-based operational and compliance process
- Create, manage, govern, and deliver trusted information
- Optimize business performance and organizational knowledge

Now there is a fifth entry point:

- Create an approach for transforming information into a trusted strategic asset that can be leveraged across applications, processes, and decisions for sustained competitive advantage

The IBM Information Agenda™ (Figure 4) is a method for transforming information into trusted strategic assets. These assets can be leveraged across applications, processes, and decisions to create a sustained competitive advantage.

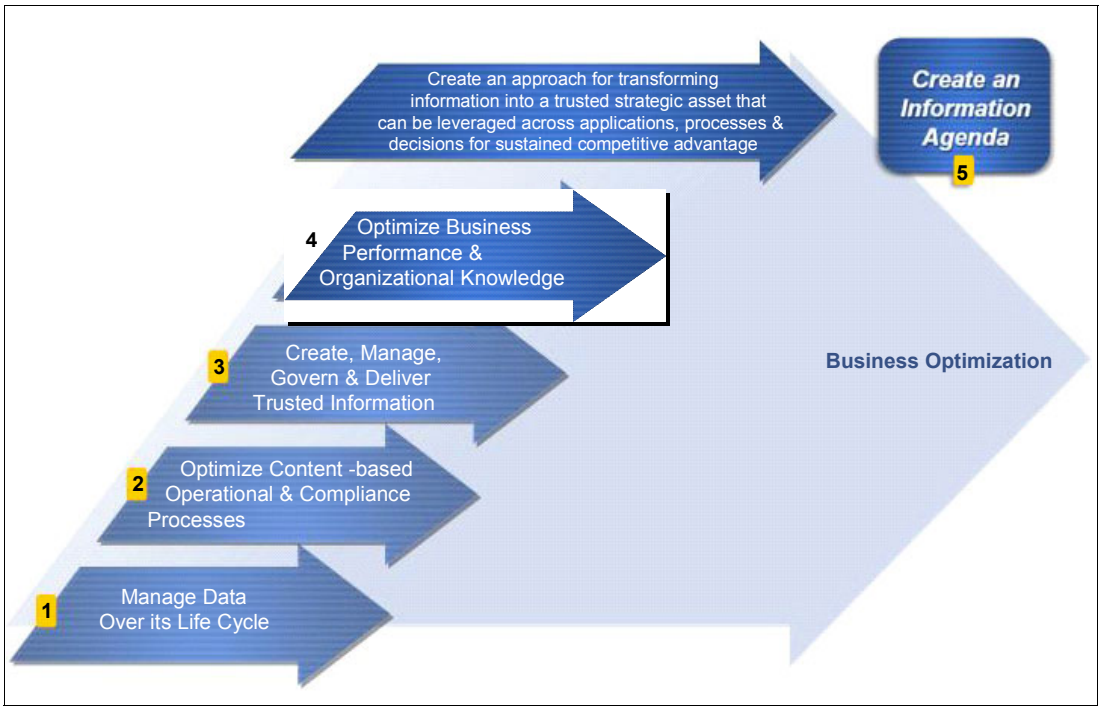

Figure 4 Creating an Information Agenda

The IBM Information Agenda integrates strategy, information governance, and enterprise information infrastructure with a comprehensive implementation roadmap. The approach is based on unique software capabilities, best practices, and deep industry knowledge. The IBM Information Agenda approach has a proven track record of helping businesses to access and share data. It can help your business become more competitive and productive by helping you to develop a plan for transforming your data into a trusted, strategic asset.

The IBM Information Agenda provides the following benefits:

- Connecting data, people and processes
- Aligning IT and business goals
- Leveraging industry-specific assets and solutions
- Establishing competency centers

The IBM Information Agenda approach has a proven track record of helping companies to respond and adapt quickly to unpredictable changes in information. IBM software and consulting services are designed to help your business develop a customized implementation roadmap in a matter of weeks, and reduce IT spending by leveraging existing investments.

The IBM Information Agenda consists of the following components:

Information infrastructure

DB2 creates the foundation for your information infrastructure, and works with IMS™ and Informix®. DB2 runs on many operating systems, such as z/OS, IBM i5, Linux, UNIX®, Windows, and Solaris. Around the information management systems is a structure that includes tools for analysis, data replication, warehouse management, content management, and information integration. Complementing the tools are key database technologies, such as XML, Service Oriented Architecture (SOA), web services, and groups of developer

communities that IBM works with to complete business solutions.

Enterprise management Products such as the IBM Information Management tools collection offer organizations a broad range of tools for everything from database management to performance analysis. The DB2 Control Center also provides tools for managing your environment. In addition, many IBM products support Tivoli® tools, which help organizations to manage enterprise information.

Business information services Business information services satisfy the major business needs of the organization. These services include Master Data Management and Entity Analytics. In addition to these IBM products, your organization can acquire applications from various independent software vendors.

Business partners IBM works with a large number of vendors and places great importance on relationships with business partners that develop and support core applications for their customers. These applications provide vital business functions, such as Customer Relationship Management and Supply Chain Management.

The need for an information architecture

An information architecture is needed in which solutions are designed that ensure that:

- The required information is collected in the first place.
- The information is up-to-date.
- The redundancy of data is kept to a minimum.
- The data is kept in a safe place.
- The data is backed up and migrated according to the business and regulatory requirements.
- The integrity of data is guaranteed.
- The access to data is performed in the most efficient manner possible.

To achieve these objectives, an information architecture is needed that focuses not only on database solutions, but also on solutions for business intelligence, business analytics, predictive analytics, data warehousing, content management, and master data management. Supporting processes such as data cleansing, data integration and replication need to be part of the overall information architecture as well.

Figure 5 illustrates a generic information architecture implemented on IBM System z and using IBM information management software solutions for all of these areas. The following layers are covered in this architecture:

Infrastructure This is the base infrastructure of the solution. It provides the platform that will hold all processes, providing qualities of service such as availability, scalability performance, security, and reliability.

Data repository The data repository layer includes the data stores for the operational data needed in traditional processing (OLTP and batch) and data in the data warehousing (OLAP) environment.

Data integration The data integration layer is needed to combine fairly unstructured data into structured information. It also helps to improve accessibility of data stores and it is here where changes in (raw) data sources are

propagated to the data warehouse and where Extract, Transform, Load (ETL) processes are executed.

Analytics and access The analytics and access layer is where information is exploited and used. Business intelligence and analytics solutions are implemented in this layer and it is here where cubing services, reports, dashboards, data access, and so on, are implemented.

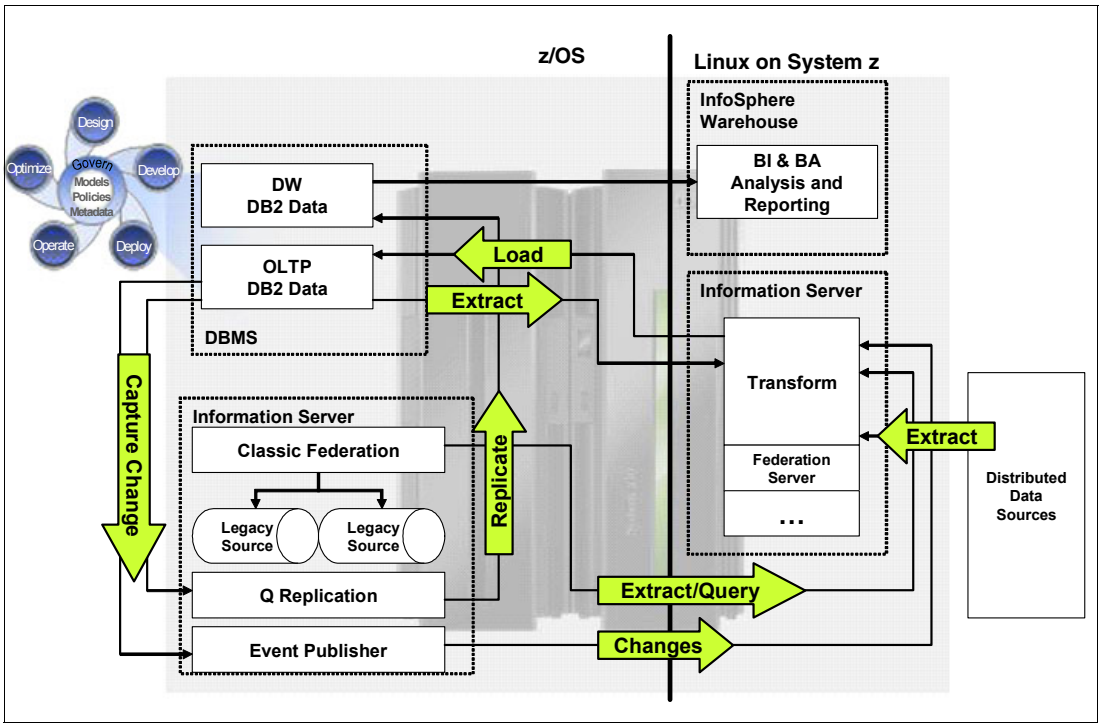

Figure 5 Information architecture on IBM System z

Strengths of System z for information management

As explained in “The ultimate challenge: coping with large quantities of data and information” on page 5, it is important to choose the right infrastructure for the current and future information management needs. In this chapter we describe the unique value proposition of the IBM System z platform, combined with the availability of a broad software portfolio for all layers of the information architecture.

The value proposition of IBM System z for information management is a combination of strengths provided by the System z hardware and z/OS operating system, the specific strengths of DB2 on z/OS, and the broad functionality of the IBM information management software solutions available.

The following topics are discussed:

- Brief history and positioning of System z
- Strengths of System z for information management
  - Green IT
  - Availability and scalability
  - Workload management
  - Reduced TCO with zIIP specialty engine
  - and more
- Strengths of DB2 on z/OS

Brief history and positioning of System z

The history of System z begins in the 1960s when computer manufacturers began to standardize hardware and software. At that time systems were dedicated to either commercial workloads or scientific workloads. With the introduction of the IBM System/360, things began to change: it was the first general purpose hardware and computer. The revolutionary S/360 was able to perform both types of computing, as long as the client, a software company, or a consultant provided the programs to do so.In fact, the name S/360 refers to the architecture's wide scope: 360 degrees to cover the entire circle of possible uses.

In 1967, IBM introduced virtualization on the mainframe. Virtualization is currently a “hot topic” but after more than four decades, no other platform has been able to achieve the mainframe's virtualization capabilities. Today, the key attribute System z delivers that differentiates it from other environments is the ability to share, at a highly granular level, the key resources of real processors, memory and I/O channels between virtual environments. As a result, System z can drive exceptionally high utilization rates and enable tens or even hundreds of workloads to run on the same single server simultaneously, while maintaining user response times.

In the early 1990s the client-server model of computing, with its distributed nodes of less-powerful computers, emerged to challenge the dominance of mainframe computers. In response, IBM designed new mainframe computers to meet the demands of changing times and a growing list of user requirements. With the expanded functions and added tiers of data processing capabilities such as web serving, autonomics, disaster recovery, and grid computing, the mainframe computer is poised to ride the next wave of growth in the IT industry.

Figure 6 illustrates the evolution of the IBM mainframe.

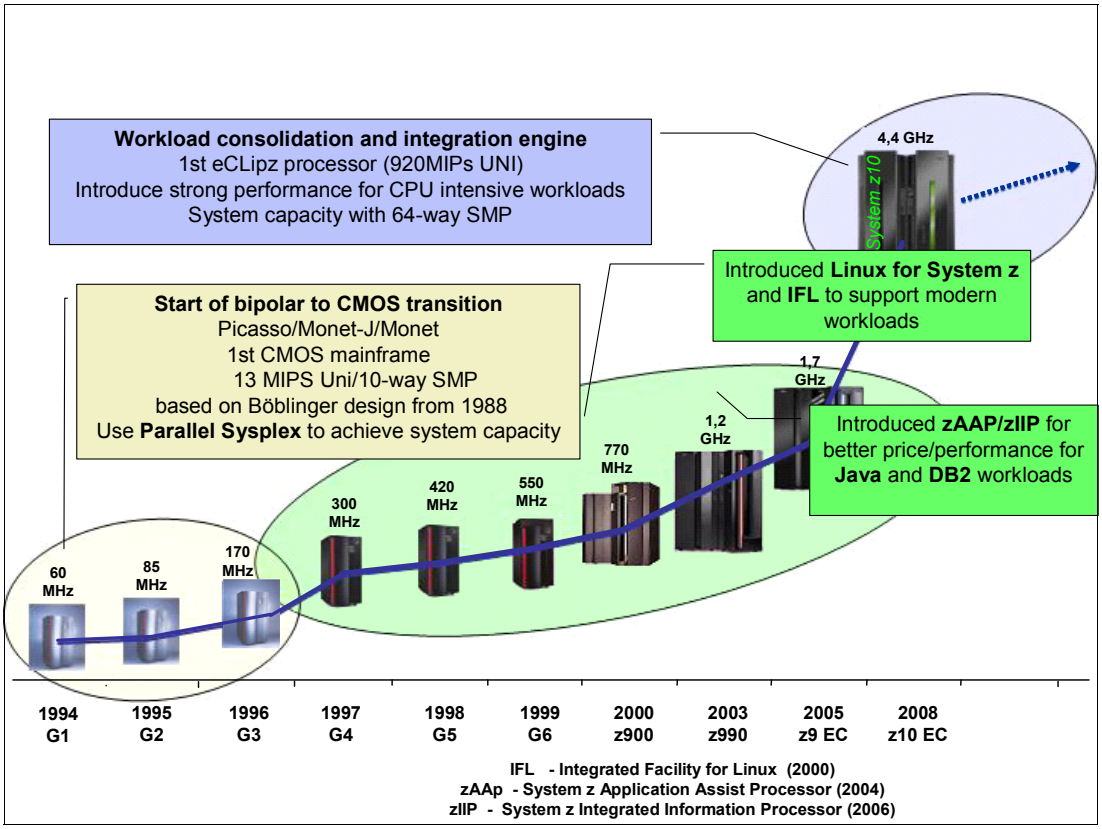

Figure 6 IBM CMOS mainframe evolution

What is next

The information in this section is based on IBM strategy and does not necessarily reflect technology that is available at the time of writing this document or will be available in the future.

The next generation of the IBM mainframe (Figure 7) will provide a new heterogeneous architecture exploiting network-attached “accelerators”. These accelerators, possibly running on IBM System x or IBM Power Blades, remain managed by an IBM System z server. One of these accelerators is the IBM Smart Analytics Optimizer; see “Exploring the IBM Smart Analytics Optimizer for DB2 for z/OS V1.1” on page 37.

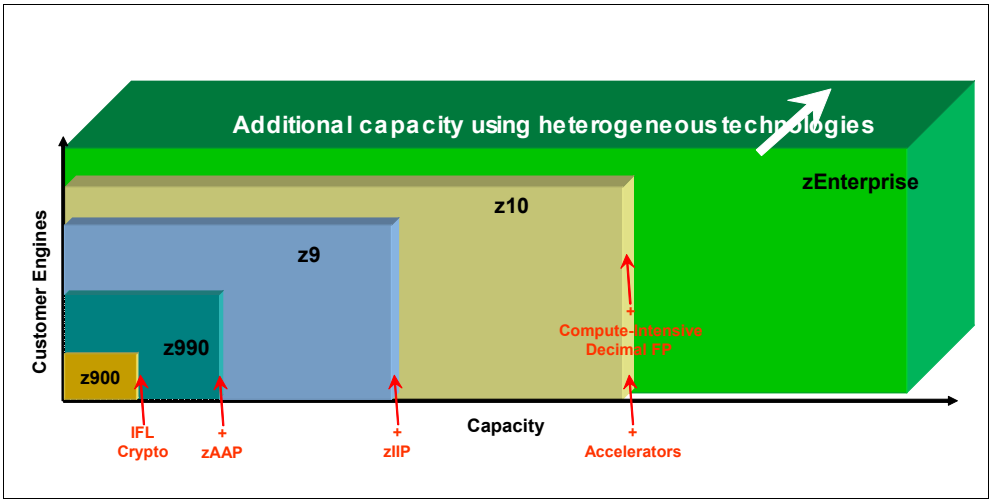

Figure 7 Strategy of the IBM mainframe

Strengths of System z for information management

System z brings the following strengths to meet your information management needs.

Green IT

In the past, huge data centers were required to house mainframes, but today the game has changed. In the 1990s, mainframe processors and most of their I/O devices became physically smaller, while their capabilities continued to grow. The most significant change occurred when IBM totally changed the core mainframe technology from bipolar to the new CMOS technology. The resulting impact on energy requirements was huge, with reductions of over 90%.

This savings is particularly meaningful when you realize that clients are spending more than twice as much on power and cooling as they are on total server purchases, and that energy costs are increasing at a rate of 2.8% per year. The power and cooling issues that data center managers face are not stand-alone challenges; they have a cascading impact on other facility issues such as wiring, floor space, and lighting.

Since that change in technology, IBM has continued to focus on improving energy efficiency, with the System z10® Enterprise Class (z10 EC) delivering approximately 15 times more performance per kilowatt hour than their original CMOS processors. Today's mainframe is highly energy efficient with unprecedented levels of performance and the ability to replace tens or hundreds of distributed servers with a single, energy-efficient System z server.

IBM has undertaken a large consolidation project, Project Big Green, to consolidate approximately 3900 distributed servers into roughly 30 mainframes, using z/VM® and Linux on System z. This project achieved reductions of more than 80% in the use of space and energy.

Availability and scalability

As businesses grow, their data processing needs also grow. Business ventures, such as mergers, acquisitions, new services, or new government regulations, can accelerate how quickly the data processing needs of the business grow. As rapid growth occurs, companies need a way to scale their business successfully.

Access to the data warehouse in a BI solution is moving from the back office to the front office, making the BI solution a business-critical application that requires the same kind of availability and security as OLTP systems. Frequently the data warehouse is large and requires careful management to ensure that business information is available when needed and that access to that data is responsive but safe and protected. The System z platform and DB2 9 for z/OS provide an excellent infrastructure for these requirements, thereby making System z the platform of choice for critical processing

The System z hardware, z/OS operating system, and DB2 for z/OS are designed with reliability characteristics such as 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.

Parallel Sysplex® clustering technology and DB2 data sharing are the answer to availability and scalability. A *sysplex* is a group of z/OS systems that communicate and cooperate with one another by using specialized hardware and software. The systems are connected and synchronized through a Sysplex Timer® or System z Server Time Protocol (STP), and Enterprise Systems Connection (ESCON®) or Fibre Channel Connection (FICON®). A Parallel Sysplex is a sysplex that uses one or more coupling facilities (CFs), which provide high-speed caching, list processing, and lock processing for any applications on the sysplex.

As a result, when applications are “enabled” for this implementation, the complete benefit of the Parallel Sysplex technology is made available. Work requests that are associated with a single workload, such as business transactions or database queries, can:

- Dynamically be balanced across systems with high performance
- Improve availability for both planned and unplanned outages
- Provide for system or application rolling-maintenance
- Offer scalable workload growth both vertically and horizontally
- View multiple-system environments as a single logical resource

The Parallel Sysplex can grow incrementally without sacrificing performance. The Parallel Sysplex architecture is designed to integrate up to 32 systems in one cluster. In a shared-disk cluster, each system is a member of the cluster and has access to shared data.

For information about Parallel Sysplex technology and benefits, refer to the Business Resiliency page at the following address:

http://www.ibm.com/systems/z/resiliency/parsys.html

Workload management

One of the strengths of the IBM System z platform and the z/OS operating system is the ability to run multiple workloads at the same time within one z/OS image or across multiple images. The function that makes this possible is dynamic workload management, which is implemented in the Workload Manager component of the z/OS operating system.

The Workload Manager (WLM) component of z/OS has proven its worth in many studies, demonstrations, and everyday work being performed in systems around the world by its ability to maximize the use of available resources. WLM has the ability to manage widely varying workloads efficiently and effectively. In doing so, you can fully use the system resources that you have available. This means, for example, that you can run your data warehouse workload together with the transactional (OLTP) workload on the same DB2 subsystem or on other DB2 subsystems on the same system (Figure 8).

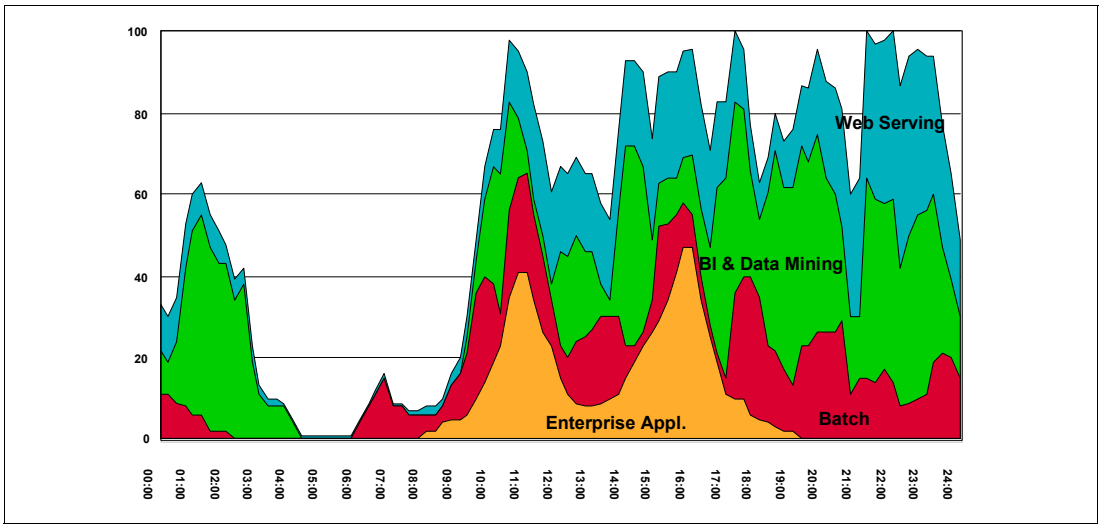

Figure 8 Running separate workloads on System z

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 closely 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 put to use whenever it becomes available.

WLM includes the following functionality:

Protect the work WLM offers several features to protect specific work from being affected by the behavior of other work. These features apply to resources such as CPU, storage, and initiator controls.

Manage server address spaces This function takes advantage of the new business unit of work (enclave) and allows WLM to manage the server address spaces of a work manager subsystem.

Routing support Routing support can provide multiple benefits such as high availability through redundancy; high performance through increased cluster efficiency, business goals, and importance considerations; and distribution of work across the Parallel Sysplex cluster according to system load.

Intelligent resource director LPAR CPU Management, LPAR weight management, Dynamic Channel Path Management (ESCON/FICON channels with Directors only), channel subsystem priority queuing.

Specialty processors	Manage and distribute the workload in each speciality processor.

For more information about data sharing and WLM on z/OS, refer to the following IBM Redpaper:

- *Workload Management for DB2 Data Warehouse*, REDP-3927

Reduced TCO with zIIP specialty engine

Reduction of total cost of ownership (TCO) can be achieved using the System z specialty engines. Running workloads on IBM System z specialty engines results in a lower hardware and software cost as a result of special reduced pricing of the specialty engines and license fees for software used for workloads running on these specialty engines.

Currently, the following specialty engines are available:

- Internal Coupling Facility (ICF) helps companies to implement a data sharing environment in the same machine.
- Integrated Facility for Linux (IFL) helps companies to consolidate workloads onto System z in a virtualized Linux environment.
- System z Application Assist Processor (zAAP) is designed to offload Java™ and some XML workload. The zAAP helps to significantly reduce the cost of running Java workloads on the IBM mainframe.
- System z Integrated Information Processor (zIIP) is designed to offload DB2 remote access, parallel queries, certain utilities, XML processing, and IPSEC encryption.

System z Integrated Information Processor (zIIP)

The IBM System z Integrated Information Processor (zIIP) is designed to maximize resource optimization. It was introduced in 2006 and DB2 immediately took advantage of it for z/OS V8. The zIIP is priced less than general purpose processors, and the millions of instructions per second (MIPs) it provides do not count toward the IBM software license fees.

In DB2 V8 and DB2 9 for z/OS, distributed requests over TCP/IP (DRDA®), query requests that use parallelism, and utilities that maintain index structures can take advantage of zIIPs. Business Intelligence application costs can directly benefit from DB2 and zIIPs. In DB2 9 for z/OS, remote (Distributed Relational Database Architecture™ or DRDA) native SQL procedures also use zIIPs and more are likely to follow. Figure 9 shows a typical data warehouse workload that can be offloaded to the zIIP.

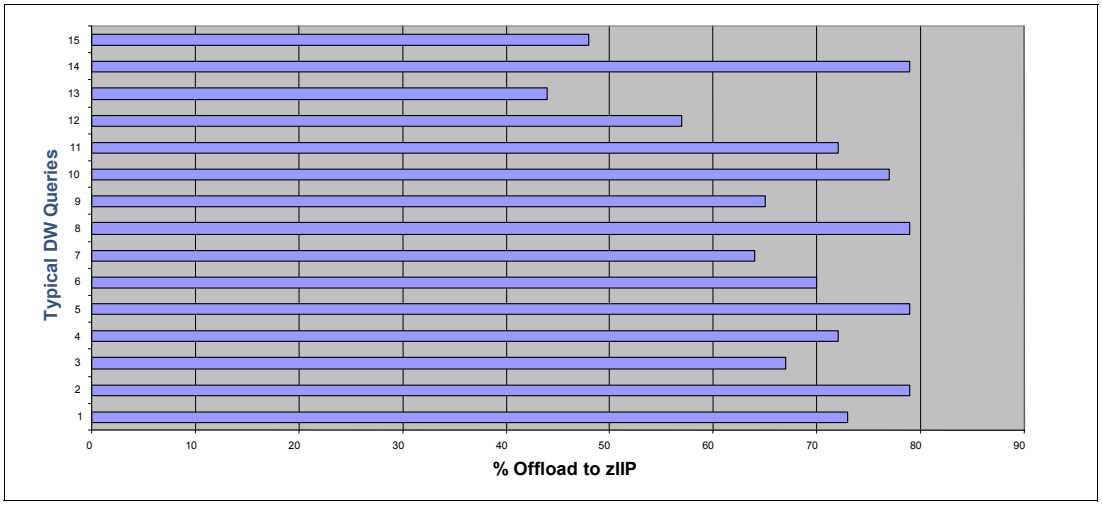

Figure 9 Example of a data warehousing workload running on a zIIP

The graph clearly shows there is a wide range of query work that can be redirected to a zIIP. The queries shown are selected TPC-H like queries executed on a 1TB warehouse. The primary work for a data warehouse is the parallel processing. Note the following points:

- For large queries that are processed in parallel, as much as 80% of the work can be redirected to a zIIP, representing most, if not all, of the costly BI queries.
- Some cases will be smaller, because part of the work will not be run in parallel.
- Small queries and the initial part of large queries (under 100 ms of CPU time) do not use the zIIP.
- If the work comes in remotely over TCP/IP and DRDA (not stored procedures), then the DRDA redirect can be used as well for the initial work.

Strengths of DB2 on z/OS

Note the following facts about DB2 for z/OS: it is used by the top 59 banks in the world, as well as by 23 of the top 25 US retailers and by 9 of the top 10 global life and health insurance providers.

DB2 for z/OS delivers high performance.

- It delivered the largest banking benchmark ever at a large bank in Asia, a record 9,445 transactions per second.
- It processed 15,000 transactions per second, almost 300,000 SQL/sec for a large Asian bank benchmark
- It supports the world's largest known peak database workload: 1.1 Billion SQL statements per hour at UPS.
- It supports the world's largest known transaction processing database: 23.1 TB at a public records agency.

DB2 for z/OS is the leading enterprise data server, designed and tightly integrated with the IBM System z mainframe to use the strengths of System z, and to reduce TCO through process enhancements and productivity improvements for database administrators and application developers. New structures, such as the ability to make changes to data definitions without disrupting online performance, continue to enhance availability and

scalability in DB2, and bottlenecks have been removed to ensure the position of DB2 as a performance leader.

The following examples show the deep synergy between DB2 on z/OS and the System z platform:

- Data sharing (to provide availability and scale out)
- Hardware data compression
- System z integrated Information Processor (zIIP) specialty engines
- Unicode conversion
- Encrypted TCP/IP communication (SSL), encrypted data
- Cross-memory, memory protection keys
- Sorting
- Multi-core, large N-way
- 1 MB page size
- Decimal float arithmetic (z10)
- 64-bit addressing and large memory
- z/OS Workload Manager
- z/OS Security Server (RACF®)
- z/OS RRS integrated commit coordinator

The following sections explain in more detail the four key strengths of DB2 on z/OS:

- DB2 data sharing
- Compression
- Partitioning
- Parallelism

DB2 data sharing

DB2 data sharing (Figure 10) improves the availability of DB2 data, extends the processing capacity of the system, provides more flexible ways to configure the environment, and increases transaction rates. You are not required to change the SQL in your applications to use DB2 data sharing. All members of a data sharing group share the same DB2 catalog and directory, and all members must reside in the same Parallel Sysplex.

The DB2 data sharing design gives businesses the ability to add new DB2 subsystems into a data sharing group, or cluster, as the need arises and without disruption. It provides the ability to perform rolling upgrades of service or versions of the software stack without any application outage.

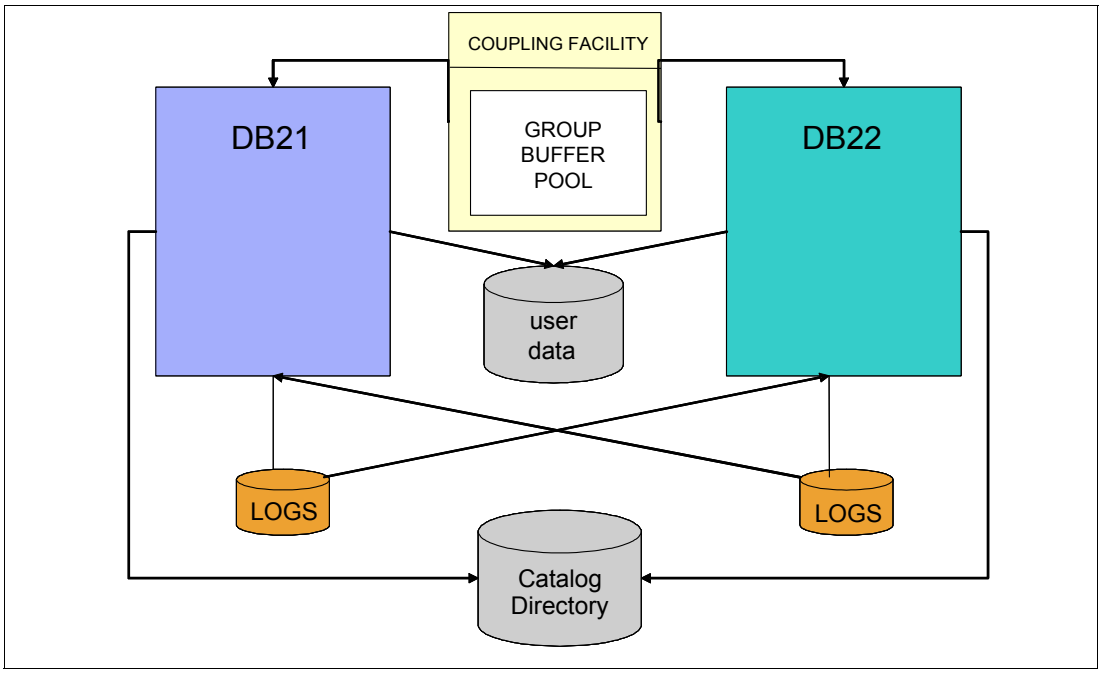

Figure 10 DB2 data sharing

With DB2 data sharing, you obtain the following benefits:

- Support for incremental growth

  A Parallel Sysplex can grow incrementally, allowing you to add processing power in granular units and in a nondisruptive manner. The coupling technology of Parallel Sysplex, along with the additional CPU power, results in more throughput for user applications. You no longer need to manage multiple copies of data, and all members of the data sharing group share a single DB2 catalog and directory.

- Workload balancing

  DB2 data sharing provides workload balancing so that when the workload increases or you add a new member to the group, you do not need to distribute your data or rewrite your applications. DB2 data sharing is unlike the partitioned-data approach to parallelism (sometimes called shared-nothing architecture), in which a one-to-one relationship exists between a database management system (DBMS) and a segment of data. When you add a new DB2 subsystem onto another central processor complex (CPC) in a data sharing environment, applications can access the same data through the new member just as easily as through any of the existing members.

  DB2 works closely with the Workload Manager (WLM) component of z/OS to ensure that incoming requests are optimally balanced across the members of a data sharing group. All members of the data sharing group have the same concurrent and direct read-write access to the data.

- Capacity when you need it

  A data sharing configuration can handle peak work loads (such as end-of-quarter processing) well. You can have data sharing members in reserve, bring them online to handle peak loads, and then stop them when the peak passes.

Compression

Counter-intuitively, compressing data can reduce elapsed time of most data warehouse-type queries. DB2 for z/OS compresses the rows on a page, so that each data page is full of compressed rows. It uses the hardware instruction along with a data dictionary to provide the most efficient compression available. The compressed data can also be encrypted, thereby saving space and implementing security requirements at the same time.

With a rule of thumb of a 50% compression rate, a compressed page contains twice the rows that an uncompressed page contains. This means that each I/O retrieves twice as much compressed data as it retrieves if the data is uncompressed. The data remains compressed in the buffer pool. This means that DB2 for z/OS can cache twice as much compressed data in its buffer pool as it retrieves if the data is uncompressed. 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.

Not all data on a compressed page is decompressed; only the row or rows that are needed by the application are decompressed. Combined with the use of a hardware instruction to perform the decompression, this decompression serves to limit the amount of additional CPU that is needed to access compressed data.

The larger amount of data retrieved in each I/O is compounded with the DB2 9 for z/OS increased prefetch quantities. This provides significant elapsed time reductions for all types of sequential processes, including the typical business intelligence queries that use table scans and index range scans. This includes sequential processes for utility access, providing benefits in terms of faster reorganizations, faster unloads, and faster recovery.

Partitioning

In a data warehouse environment, the size, growth estimates, or both of certain tables, such as fact tables and history tables, pose a challenge in designing and administering the table space. They also present challenges in ensuring query performance, and loading and deleting data. Partitioned table spaces, as they have evolved with the implementation in DB2 9 for z/OS, help resolve most of these problems.

Your data can be partitioned as follows:

Range partitioning table space — This is a universal table space or non-universal table space partition table space that uses ranges to define the various partitions. The number of partitions and the range of each partition is user-controlled.

Partitioned by growth table space — This table space is always a universal table space. Additional partitions are added to a partitioned by growth universal table space as additional space is needed until a predefined maximum number of partitions is reached. Partitioned by growth universal table space ranges and new partitions added to a partitioned by growth universal table space are completely managed by DB2. A user has no control over the ranges used or when a partition is added.

Table partitioned — Partitioning keys are define within the table's description at CREATE TABLE time. Table partitioning does not use a partitioning index. It also separates clustering from partitioning.

Data partitioned secondary index This index is a type of partitioned index available in DB2 for z/OS Version 8 and later. A data partitioned secondary index has the same number of index partitions as the table space has partitions. It is defined with columns other than the columns that are used to define table-controlled partitions.

Parallelism

You can significantly reduce the response time for data- or processor-intensive queries by taking advantage of the ability of DB2 to initiate multiple parallel operations when it accesses data in a data warehouse environment.

Note the following types of parallelism.

- Query I/O parallelism manages concurrent I/O requests for a single query.
- Query CP parallelism enables true multitasking within a query.
- Sysplex query parallelism enables a large query across separate DB2 members in a data sharing group.
- DB2 can use parallel operations for processing the following types of operations:
  - Static and dynamic queries
  - Local and remote data access
  - Queries using single table scans and multi-table joins
  - Access through an index, by table space scan or by list prefetch
  - Sort
- Utilities parallelism:
  - Unload & Load
  - Reorg
  - Copy
  - and so on

Why to use DB2 on z/OS

Why use DB2 for z/OS? The following reasons show why it is the best choice to implement the data warehouse:

- Perhaps the majority of source systems are on z/OS within IMS, VSAM, DB2, or sequential files, or there is a requirement for tight integration with existing resources and systems on the System z platform.
- The operational data store already exists on System z as several of the data warehouses and data marts.
- Existing skills and investments are on the System z platform.
- The company already maintains a System z-centric IT solution due to its favorable cost of ownership and comfort.
- The company is implementing an operational BI application with embedded analytics within applications. These types of applications can use the System z transaction scalability capabilities.
- There is a requirement for a true real-time operational data store:
  - Operational data is already on the System z platform.
  - Data must be virtually in sync with the operational data.
  - Availability, security, and resiliency requirements are high.
  - Auditable data warehouse requirements exist.

- Independent software vendor (ISV) packages, such as SAP and PeopleSoft on System z, offer both transactional (OLTP) and informational (warehouse and BI) systems. System z supports multiple images of SAP and other solutions, which simplifies support for these complex applications.

  These packaged applications, which have tightly integrated components, have always made it desirable for the operational data and the warehouse to be housed in the same environment. Collocation reduces operational complexity, allowing for the reuse of skills and infrastructure in the form of processes, tools, and procedures.
- A desire already exists to consolidate distributed marts or data warehouses to an existing System z data serving platform. The customer possibly has spare System z capacity.

IBM InfoSphere on System z

Organizations face an information challenge that begins with locating information, understanding what it means, and getting the information when it is needed in the form needed. Compounding difficulties include the uncertainties of whether the information can be trusted, and how to control it. The challenges continue to grow if businesses cannot ensure that they have access to authoritative, consistent, timely, and complete information.

IBM delivers market-leading solutions to critical information-intensive business problems, allowing clients to achieve new levels of innovation-building trusted information through best of breed information integration and data warehousing and MDM capabilities.

There are four major parts that describe the IBM InfoSphere™ solution:

- IBM InfoSphere Foundation Tools and IBM Industry Models for design and governance
- IBM InfoSphere Information Server for System z for information integration
- IBM InfoSphere Warehouse for data warehousing
- IBM InfoSphere Master Data Management for master data management

This chapter discusses these areas in more detail; Figure 11 provides an overview.

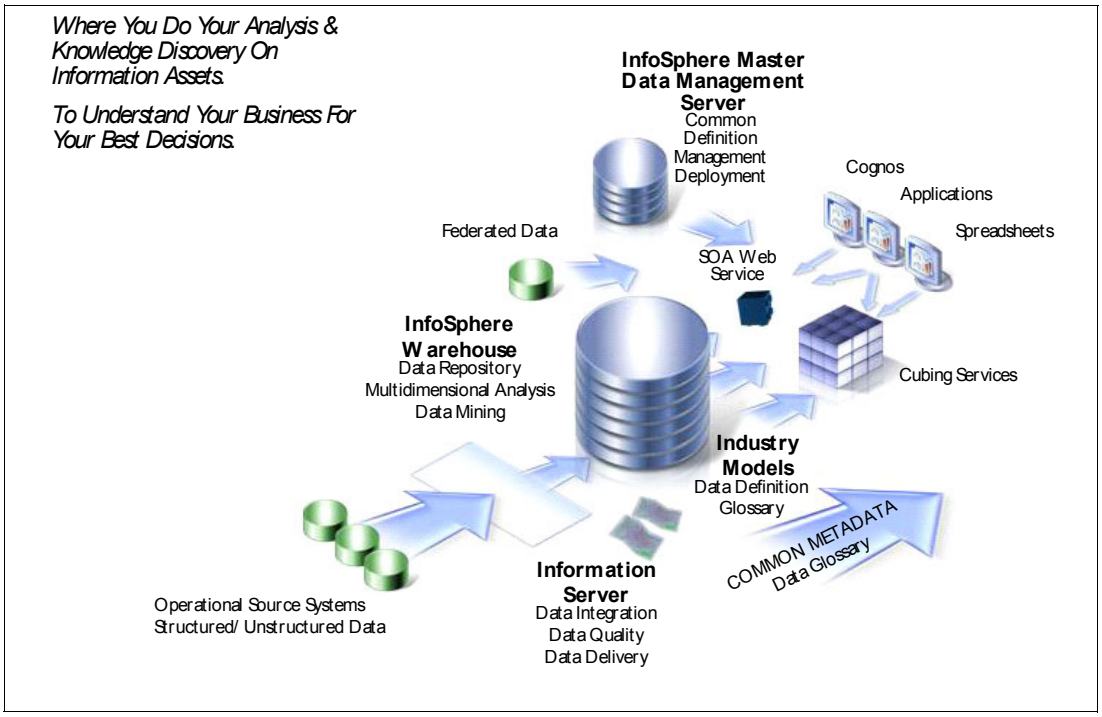

Figure 11 InfoSphere offering end-to-end - The IOD value chain

InfoSphere Foundation Tools and Industry Models

With data spread all over your systems, to build trusted information you need to understand how this data is related, how new data is to be designed, and how to govern this data through its life cycle. InfoSphere Foundation Tools helps you to discover, design, and govern your data. It can help you with the following tasks:

- Shorten development cycle times
- Increase operational efficiency
- Eliminate and streamline duplicate processes
- Promote team collaboration
- Go to market quicker
- Become more responsive to client demand
- Minimize project risk

InfoSphere Foundation Tools are an integrated set of tools that help you to discover and map your information sources, understand the structure, content and quality of information available, and assess the gap between available information and what is needed for a specific project, such as an ERP implementation, a data warehousing project and so on. InfoSphere Foundation Tools also play a key role in data governance after you have implemented your project. For example, after an SAP ERP deployment, InfoSphere Foundation Tools can query your ERP deployment on a regular basis (typically, weekly) to assess the quality of information over time and provide you with valuable information such as the information quality for a specific part of your business that is deteriorating over time.

For more information, see the resource listed:

`http://www.ibm.com/software/data/infosphere/foundation-tools/index.html`

IBM Industry Models combine deep expertise and industry best practice in a usable form, called a “blueprint” by both business and IT communities, to accelerate industry solutions. Part of the IBM InfoSphere portfolio, the industry models are based on the experience of more than 500 clients, and more than 10 years of development.

The IBM Industry Models consist of two parts:

- Comprehensive data models containing data warehouse design models, business terminology model and analysis templates to accelerate the development of business intelligence applications.
- For selected industries, best practice business process models with supportive service definitions for development of a service-oriented architecture (SOA). The result is acceleration of projects and reduced risk.

The following data and process models are available:

- IBM Banking Data Warehouse
- IBM Banking Process and Service Models
- IBM Insurance Information Warehouse
- IBM Insurance Process and Service Models
- IBM Health Plan Data Model
- IBM Retail Data Warehouse
- IBM Telecommunications Data Warehouse
- IBM Financial Markets Data Warehouse
- IBM Banking and Financial Markets Data Warehouse

For more information about IBM Industry Models, see the resource listed:

`http://www.ibm.com/software/data/industry-models/`

InfoSphere Information Server for System z

Information Server for System z delivers all the functions required to integrate, enrich, and deliver information that you can trust for your key business initiatives. With rich functionality, broad connectivity to heterogeneous sources, and a unified metadata-driven approach, it provides a strong foundation for enterprise information architecture.

InfoSphere Information Server for System z is a fully integrated software platform that profiles, cleanses, transforms and delivers information from mainframe and distributed data sources alike to drive greater insight for the business without added IBM z/OS operational costs.

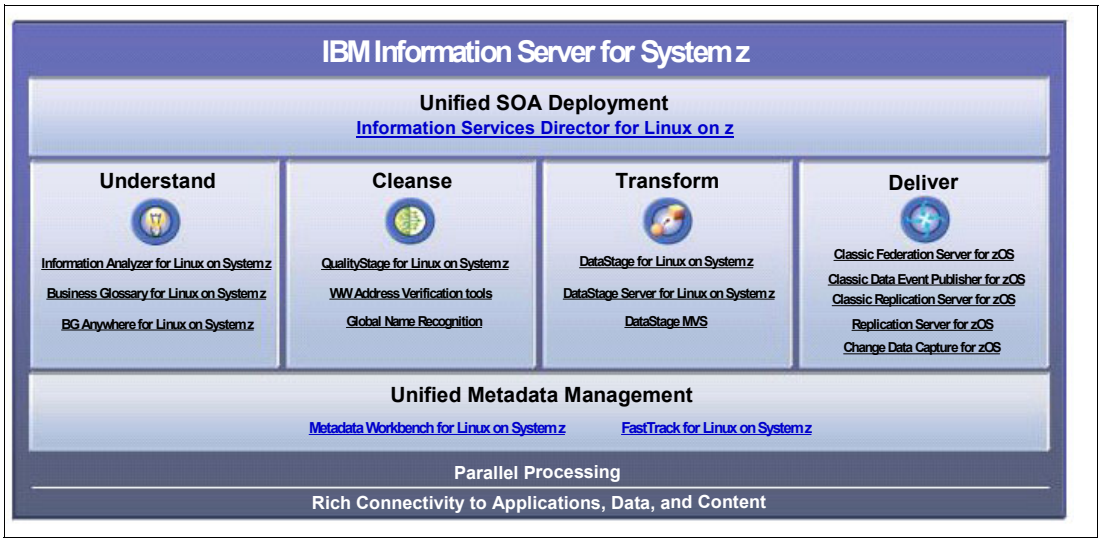

Figure 12 IBM InfoSphere Information Server

The InfoSphere Information Server solution, as illustrated in Figure 12, can be divided in four major layers:

- Understand your data
- Cleanse your information
- Transform your data into information
- Delivery your information

These all share a common layer known as *Platform Services*. This shared infrastructure ensures consistency, improves performance, and reduces the cost of administering the platform. The services include a shared parallel processing engine, allowing the processing to scale linearly as you add hardware; common connectivity to nearly any data source; shared metadata services to enable seamless end-to-end information sharing; and common administration and deployment services to reduce administration cost and complexity. Each module comes with these platform services built in, and if you add additional modules, they simply snap into the shared infrastructure and begin taking advantage of it.

The following sections explain these layers in greater detail.

Understand your data

Understand your data and the terms that describe it to ensure consistency, trust, and effective communication between the business user and technology user. IBM Information Server for System z enables discovery, definition and modeling of information content and structure.

This layer includes the following IBM solutions:

- IBM InfoSphere Discovery

  This solution provides market-leading capabilities to automate the identification and definition of data relationships across the complex, heterogeneous environments prevalent in IT today. Covering every kind of data relationship, from simple to complex, InfoSphere Discovery provides a 360° view of data assets.

- IBM InfoSphere Information Analyzer

  This solution focuses on helping you understand the structure and contents of the various source systems within your environment.

- IBM InfoSphere Business Glossary

  This solution provides a web-based tool for authoring, managing, and sharing business metadata.

Cleanse your information

Information Server for System z provides data cleansing capabilities that help ensure auditable data quality and consistency by standardizing, validating, matching and merging information to create comprehensive and authoritative information for multiple uses. Quality logic is designed visually and deployed universally, helping to ensure data consistency and accuracy across the enterprise.

This layer includes the following IBM solutions:

- IBM InfoSphere QualityStage™

  This solution standardizes, cleanses and de-duplicates data, ensuring a complete, accurate view of information across heterogeneous sources.

- IBM InfoSphere Global Name Recognition

  This solution recognizes, matches, and classifies personal and business names across cultures, and enriches name understanding with gender, origin, and structure.

Transform your data into information

This layer is used to transform and enrich information for new uses in new contexts. It repackages and repurposes information throughout the enterprise to ensure the information is provided where it is needed and in the format that is appropriate for the user or application. It transforms data from its source system-centric format into new formats more appropriate for your project goals, for example, dimensionalized for loading into a data warehouse, or restructured to load into an ERP system.

This layer includes the following IBM solution:

- IBM InfoSphere DataStage®

  This solution integrates, transforms, and delivers data on demand across multiple sources and targets including databases and enterprise applications.

Deliver your information

This layer is used to access and deliver your information when and where it is needed. Delivery can be performed through a variety of mechanisms, including bulk movement of data into new repositories, event-driven movement based on message queues or changed data events in source systems, or on-demand access through SOA or federated queries. Whether the information is delivered on demand, through federation, or on a timed or event-driven basis, it can be moved in bulk from location to location or accessed in place reusing the same core logic.

This layer includes the following IBM solutions:

- InfoSphere Change Data Capture

  This solution provides real-time change data capture for dynamic warehousing, real-time reporting, synchronization, replication, and event detection with high scalability and flexibility to implement multiple environments and topologies.

- IBM InfoSphere Federation Server

  This solution enables access and delivery of diverse and distributed information as though it were in one system.

InfoSphere Warehouse on System z

InfoSphere Warehouse on System z is part of a comprehensive IBM data warehousing and business intelligence (BI) solution for System z. It represents a highly scalable, highly resilient, lower cost infrastructure way to optimize a DB2 for z/OS data warehouse.

InfoSphere Warehouse on System z significantly improves query performance for users who want to employ OLAP techniques to drill down into specific data stored in DB2 for z/OS. The combination of the System z platform and InfoSphere Warehouse gives customers the ability to support near real-time analytics based on core business data managed in DB2 for z/OS, helping customers gain additional competitive advantage and value from their operational data.

InfoSphere Warehouse on System z (Figure 13) provides powerful features which include:

- Delivering multidimensional data with no-copy OLAP analytics
- Building the data warehouse with advanced design and physical modeling using Design Studio
- Populating the data warehouse with simplified SQL-based data movement and transformation with the SQL Warehouse tool

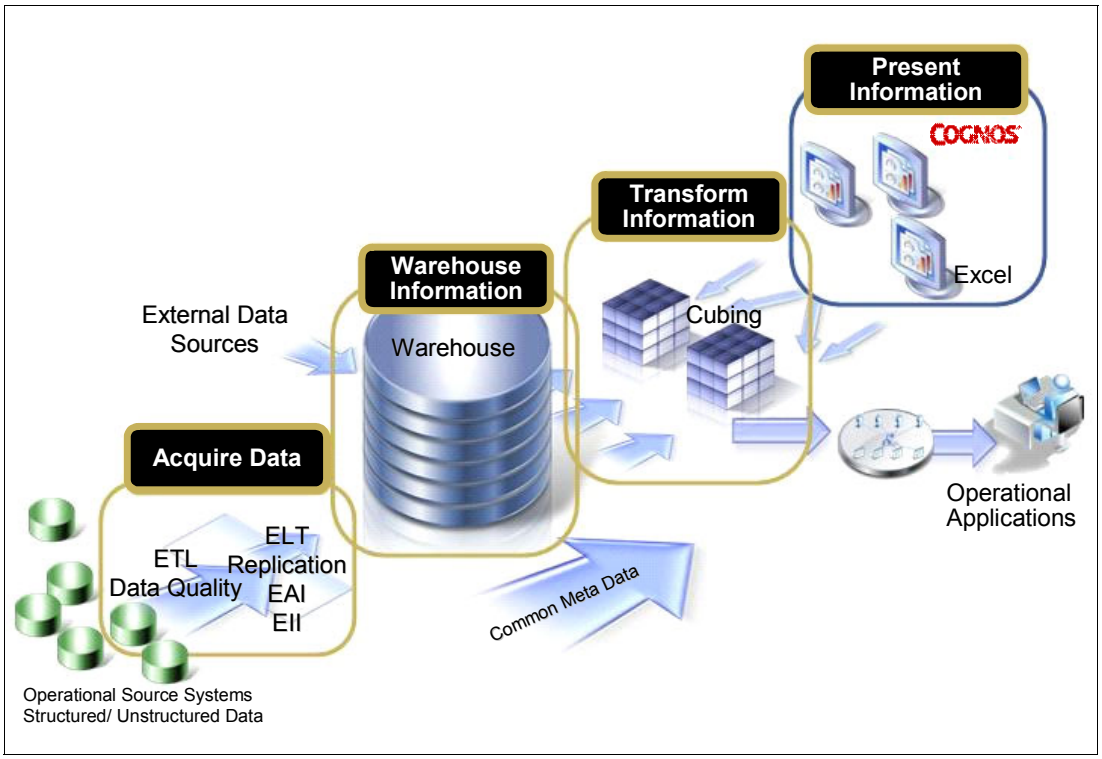

Figure 13 InfoSphere Warehouse - The integrated stack on System z

InfoSphere Warehouse on System z complements IBM DB2 for z/OS with a highly scalable, low-cost way to design, populate and optimize a DB2 for z/OS data warehouse, datamart or

operational data store for BI applications like IBM Cognos® 8 BI. It provides physical data modeling, Extract, Transform and Load (ETL), and multidimensional server functions on top of an existing DB2 for z/OS implementation and forms the core of the data warehousing environment. InfoSphere Warehouse on System z can also facilitate improved query performance through the use of Cubing Services caching for Multidimensional Expressions (MDX) query support, helping organizations save time and CPU costs.

Figure 14 shows the implementation of the InfoSphere Warehouse on System z solution.

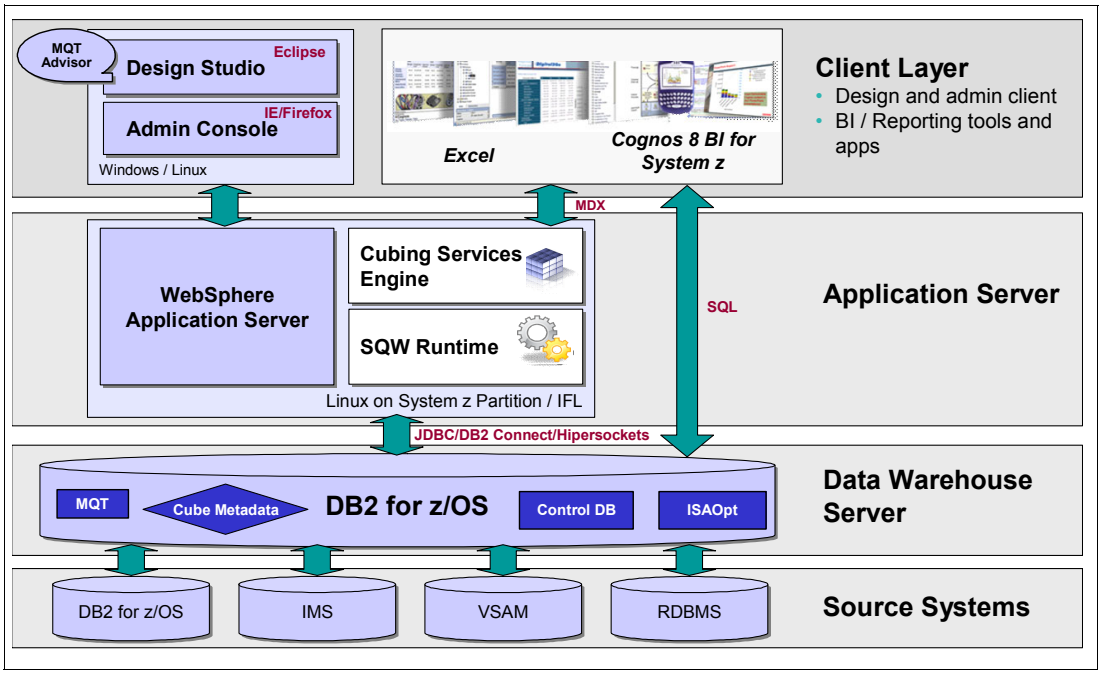

Figure 14 The InfoSphere Warehouse solution on System z

Features included in IBM InfoSphere Warehouse for DB2 for z/OS:

- Access to multidimensional data with no-copy OLAP analytics
- Eclipse-based Design Studio
- Browser-based Administration Console
- Physical Modeling from InfoSphere Data Architect
- Simplified SQL-based data movement and transformation with SQL Warehousing Tool (SQW)

InfoSphere Warehouse on System z users can also build and deploy multidimensional analytics with Cubing Services, which empowers users to create intuitive and complex ad hoc queries about their enterprise, such as a product's profitability for a defined quarter across a specific territory. Fully integrated into the InfoSphere Warehouse on System z design and administrative warehouse tools, Cubing Services provides the ability to achieve lower design and maintenance costs over external cubing solutions.

Master data management on System z

One of the key elements of an information architecture is the management of *master data*. This section summarizes the IBM solution for Master data management on System z.

What is master data

Master data consists of the key facts describing your core business entities, such as clients, partners, employees, products, bills of material, and accounts. These entities are typically used repeatedly across the entire organization and in various business processes.

Unfortunately, despite its importance master data is typically scattered within heterogeneous application silos across the enterprise. Another problem is that master data is kept in multiple redundant data stores, thus complicating any effort to maintain the state of the master data (consider the typical example of maintaining the attributes of a large bank's accounts in various applications and in many other data stores).

In general, master data can be categorized according to the kinds of questions it addresses, and three of the most common questions are "who," "what," and "how?" These three questions are reflected by the *Party*, *Product*, and *Account* domains of master data, respectively.

Each of these domains represents a class of things. For example, the Party domain can represent any kind of person or organization, including clients, suppliers, employees, citizens, distributors, and organizations. In turn, each kind of Party shares a common set of attributes, such as the name of the Party, where it is located (a party can have multiple locations such as home, work, vacation home, and so on), how to contact it, what kind of relationship the organization has with the party, and so forth.

Similarly, the Product domain can represent all kinds of things that you sell or use, from tangible consumer goods to service products such as mortgages, telephone services, or insurance policies. The Account domain describes how a Party is related to a product or service that the organization offers. What are the relations of the Party or Parties to this account, and who owns the Account? Which Accounts are used for which Products? What are the terms and conditions associated with the Products and the Accounts? And how are Products bundled?

Why you need to manage master data

Organizations need to gain control over master data to meet fundamental strategic objectives such as growth, revenue generation and cost reduction. The concept of controlling master data is called *master data management*.

Master data management is a strategy, not a product. The fundamental principle of master data management is that master data is decoupled from operational, transactional and analytical systems into a centralized independent repository or "hub." This centralized information is then provided to service oriented architecture (SOA) business services so data is managed independently of any single line of business, system or application. This strategy enables enterprises to identify common functionality for all systems and applications, and then support the efficient, consistent use of business information and processes.

Figure 15 illustrates a simplified view of the concept of master data management.

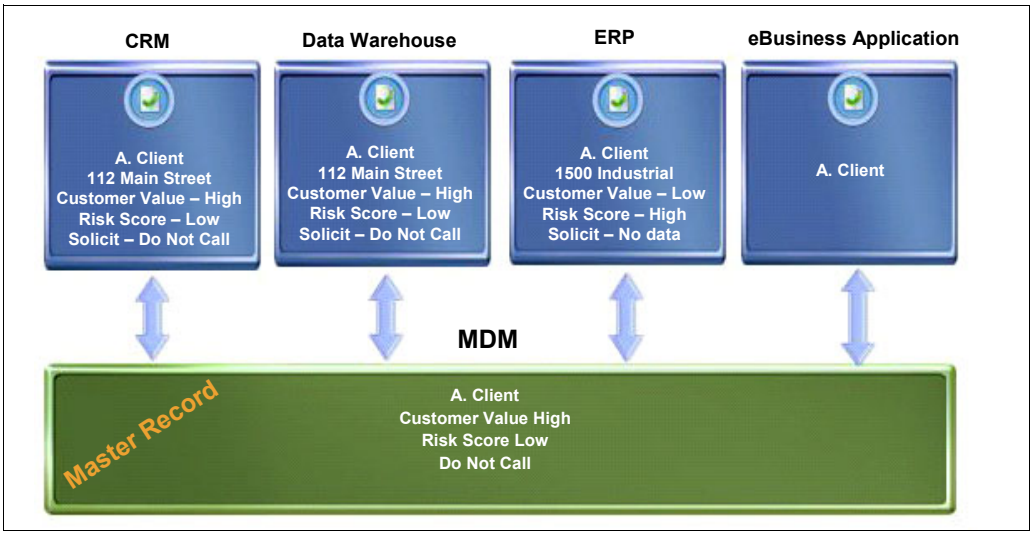

Figure 15 Master data management overview

Implementing master data management

Master data management decouples master information from individual applications, becoming a central, application-independent resource and containing configurable functionality to maintain and be the system of truth for master data. Master data management is an integration of common data functionality into an enterprise application supporting all applications through business services that create and consume master data. It simplifies ongoing integration tasks and new application development. It ensures consistent master information across transactional and analytical systems, addressing key issues proactively such as data quality and consistency, rather than “after the fact” in the data warehouse.

Master data management is complementary to application processes. It provides applications with accurate and complete data about all key business entities.

Master data is dynamic, not static; therefore a master data management application must be able to manage the complete master data life cycle. Applications that use master data play various roles, and the same application can create data, consume it, and require analysis of it. A master data management application must be able to provide mechanisms for consistent use of master data across the organization, and integrate the functional styles to manage data life cycle processes.

IBM solution for master data management

IBM InfoSphere Master Data Management allows you to gain a unified, complete, consistent and standardized view of your data to drive critical business decisions. It also helps you to gain control of your data, reduce information errors, and eliminate duplicate data, all of which will enable you to meet growth, revenue-generation, and cost-reduction goals.

The IBM MDM portfolio is anchored by two primary products: IBM InfoSphere Master Data Management Server and IBM InfoSphere Master Data Management Server for Product Information Management. Both can deliver on MDM requirements across a number of tactical needs, including two of the most common MDM projects: Customer Data Integration (CDI) and Product Information Management (PIM).

CDI projects revolve around customer data and usually focus on enabling real-time access to this information in support of business applications or processes. PIM projects generally include the technology, people, and processes required to create and define a single view of product information across the enterprise. As organizations move to the latter phases of MDM, these technologies can be deployed together to fulfill long-term strategic requirements and generate even greater value for the business.

IBM InfoSphere Master Data Management Server

IBM InfoSphere Master Data Management (MDM) Server is an industry-leading software product that provides you with the capability to manage master data involving clients, products, accounts, suppliers, members, citizens, employees, prospects, guests, agents, items, product bundles, parts, agreements, and more. IBM InfoSphere MDM Server manages master data for all business processes that consume it and is designed to be the system of record for master data.

At its core, MDM server is a master database with business services for maintaining master data. MDM Server is a real-time operational system. To prepackage a significant number of services, it is designed around a comprehensive data model. Both the services and the model are designed to accommodate configuration and change.

IBM InfoSphere Server for Product Information Management

IBM InfoSphere MDM Server for PIM provides a single, up-to-date repository of trusted product, location, supplier and other information by building business processes that capture, enrich, and maintain a 360-degree view of data. Clients benefit from its flexible data model, business process collaboration tools, data governance, aggregation and syndication capabilities and customizable user access privileges to represent their specific information needs. MDM Server for PIM helps clients accomplish a range of business objectives, including introducing new products and service bundles quickly, executing multichannel strategies, managing business-to business and business-to-consumer e-commerce sites, maximizing operational efficiencies, providing consistent client care, working effectively with partners and suppliers, communicating with employees and publishing relevant materials.

IBM InfoSphere Master Data Management Server on System z

IBM InfoSphere MDM Server is available for Linux on System z, and it benefits fully from the underlying System z virtualization capabilities. When combined with a DB2 on z/OS database to store all the master data, all of the previously described benefits of using DB2 on z/OS apply, as well. The connectivity between the MDM Server on Linux on System z and DB2 on z/OS takes place using transparent database drivers and using HiperSockets™ for fast performance.

The power of Linux on System z

The deployment of more and more distributed servers often results in a level of complexity that is unmanageable and increasingly expensive. Virtual Linux servers can help you to unify your IT infrastructure by consolidating that environment onto a single centralized server, potentially providing better resource utilization, easier maintenance, and more effective and efficient operations.

Linux on System z inherits the hardware strengths that make the mainframe a superior platform for information management: security, availability, stability, performance, reliability and so on. In addition to these strengths, Linux on System z exploits the virtualization capabilities of the platform, thereby leading to unparallelled flexibility and scalability. Running

on a special engine known as Integrated Facility for Linux (IFL), Linux on System z can help to lower your IT cost, whether for a small, medium, or large enterprise. Linux on System z allows for massive consolidation, exploiting the leading capabilities of the System z servers to run multiple and mixed mission-critical and infrastructure workloads concurrently.

Most of the solutions of InfoSphere Server runs on Linux on System z. With InfoSphere Server running on Linux for System z, you can take advantage of not only the IFL but also the System z Integrated Information Processor (zIIP) when using DB2 for z/OS as the underlying database management system. The communication between Linux on System z and z/OS uses HiperSockets, which is a cross-memory TCP/IP communication proving high performance and security.

In summary, combining Information Server processes on Linux on System z and DB2 on z/OS provides you with the following benefits:

- Scalability and flexibility by using virtualized Linux servers. The number of Linux servers can be easily increased and decreased based on workload needs, departmental needs, security requirements, and test and development needs.
- Keeping the operational data in DB2 on z/OS provides the highest Quality of Service with respect to storing and managing your precious data.
- Communication between processes running on Linux on System z and DB2 on z/OS takes place over the high-performing and secure HiperSockets, which is part of the System z hardware.
- Specialty engines (the IFL on the Linux side, and the zIIP on the z/OS side) provide lower cost.

Figure 16 shows an example with DB2 on z/OS as an operational data store and the database for InfoSphere Warehouse and various Linux on System z partitions for InfoSphere Server and Cognos BI functions.

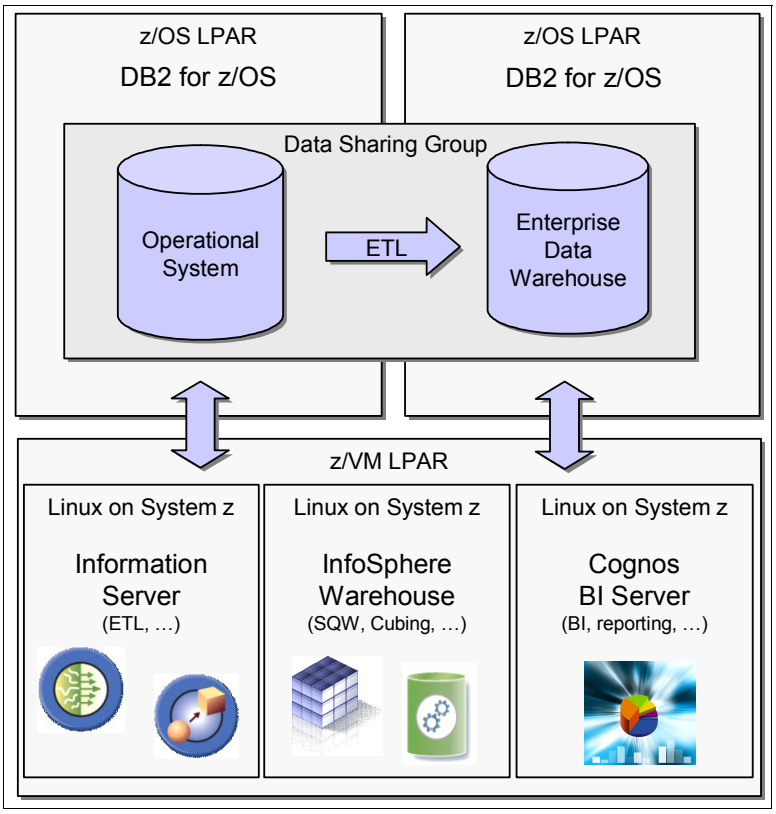

Figure 16 Example of Linux on System z and z/OS working together

IBM Cognos BI for Linux on System z

IBM Cognos BI for Linux on System z helps organizations to make better informed, faster, and more strategically aligned business decisions. It delivers the complete range of BI capabilities on a single Service Oriented Architecture (SOA). Organizations can author, share, and use BI that draws on data across all enterprise sources for better business decisions.

Note the following main features of IBM Cognos BI for Linux on System:

- A full range of BI capabilities, including reporting, analysis and dashboards, that provide all user communities with relevant information in the right structure, at the right time, and in the right place
- An open enterprise-class platform that delivers complete, consistent, and timely information across diverse user communities with cost-effective scalability
- Frameworks and proven practices to help give projects a good start, plus the expertise to enable success on the journey to performance management

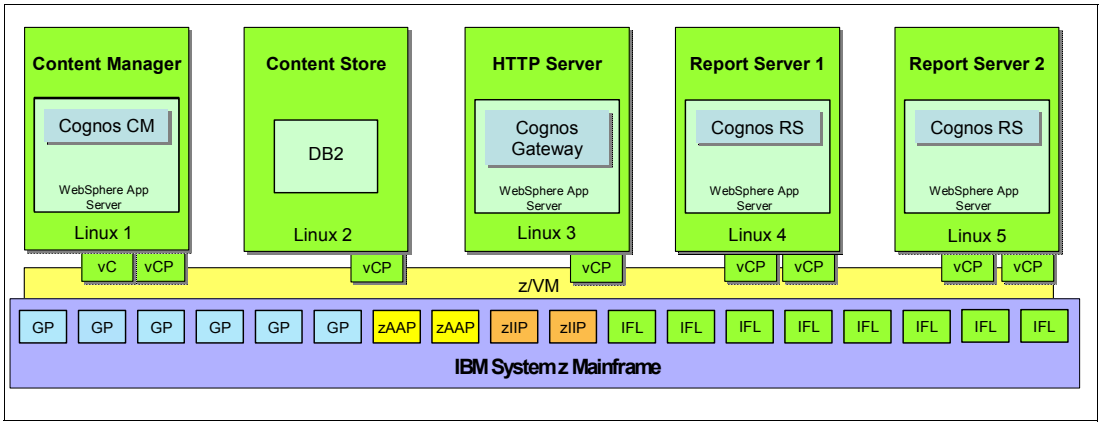

Figure 17 IBM Cognos 8 BI in a z/VM "distributed" deployment

Cognos 8 BI consolidation to Linux on System z case study

In November 2007, a test was performed at the IBM Innovation Center for Business Partners with Cognos 8 BI Version 8.3 to demonstrate consolidation to Linux for System z. The testing used a typical distribution of reporting and analysis across an organization. The goal of the testing was to prove Cognos 8 BI and the Cognos 8 platform is built for enterprise deployments.

The results discussed in this publication are specific results achieved during this test. Another test, with even the slightest difference in workload, processor types, equipment types, software levels and so on will probably produce other results. Therefore, the results in this section can only be used as an indication of what could be possible, but is not a guarantee that another test will produce the same results.

The results of this test can be summarized as follows:

- Up to 80% saving in IT cost
- Up to 96% less hardware
- 760 x86 processor cores versus 26 IFLs
- Potential for dramatic reductions in software expense for processor-based licenses
- Potential reductions in power and cooling
- Up to 93% savings in kilowatts and energy costs in this scenario
- Up to 46% less space
- Up to 89% people savings
- Increased processor utilization
- Industry-leading security

Figure 18 shows a comparison between the various platforms used in the test.

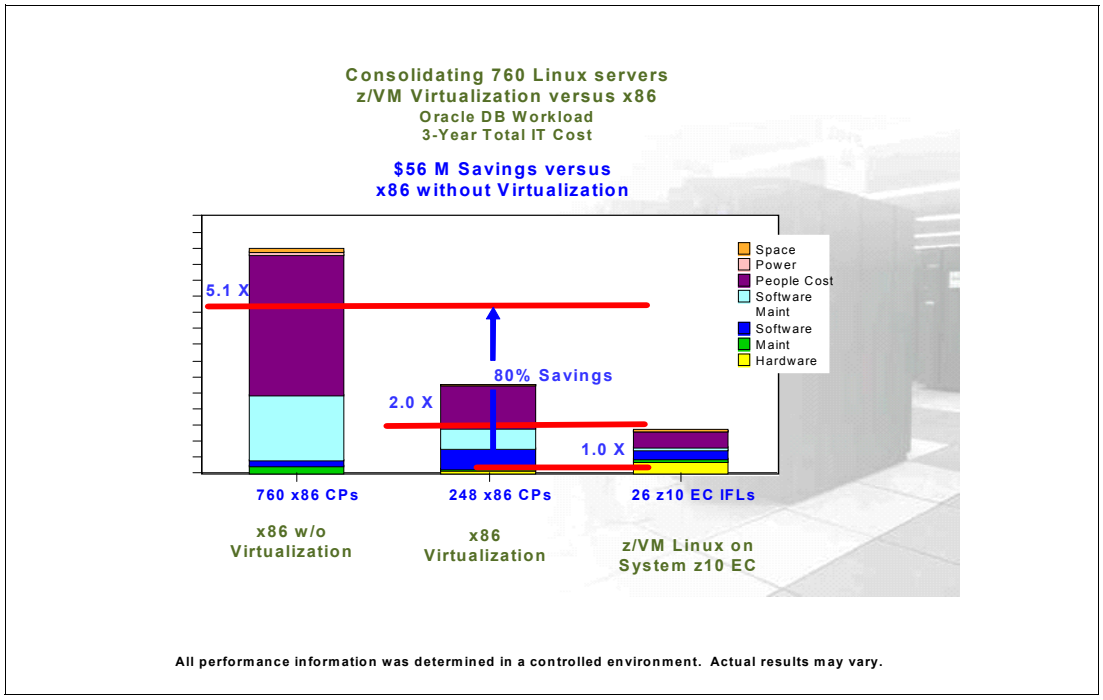

Figure 18 Consolidating Cognos 8 BI workload to Linux for System z

The scalability of the solution running in Linux for System z on a System z 10 EC has proven to be linearly scalable under a workload of up to 90000 named users, as shown in the graph in Figure 19.

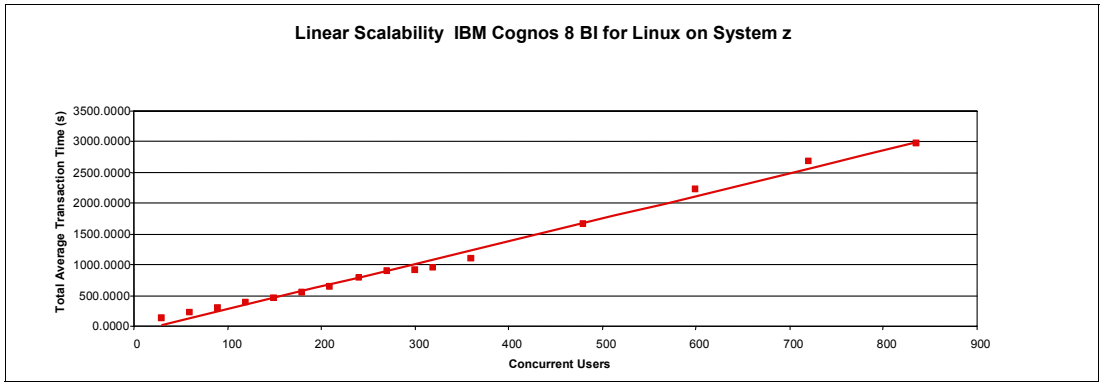

Figure 19 Cognos 8 BI scalability on a System z10 EC

Exploring the IBM Smart Analytics Optimizer for DB2 for z/OS V1.1

This chapter introduces and explores a new and powerful solution for accelerating certain queries with a magnitude. This solution is called IBM Smart Analytics Optimizer for DB2 for z/OS V1.1. It is a combination of DB2 on z/OS enhancements, the all new zEnterprise BladeCenter Extension Model Model 002 for the zEnterprise 196, a new state-of-the-art way of dealing with data marts, and tooling for development and deployment.

A solution for accelerating queries

Business intelligence (BI) and data warehousing (DW) are becoming more mainstream and more mission-critical to enterprises of all sizes. At the same time, the amount of data a business generates continues to grow at an ever-increasing rate. Even with the speed of current data growth, corporations must still make sense of it, turning data into insight and making decisions faster than ever before. Business users are finding it increasingly difficult to locate the right information within an overwhelming information explosion.

“Business intelligence has “invaded” the operational space in a big way, offering in-line analytics, real-time or near real-time decision-making support for all employees in the enterprise. Today's BI environment includes three forms of BI - strategic, tactical, and operational.” [Claudia Imhoff, 2007[1]]

Operational BI refers to the methodologies, processes, and applications used to report on and analyze operational business performance.

Traditional BI consisted mostly of executives analyzing historical data and performing strategic analysis on general questions; for example, how an organization performed against certain criteria over a given time period.

Now, operational intelligence or operational BI has moved BI into all levels of an organization, allowing not simply executives, but everyone within the organization, to incorporate proactive

[1] Operational Business Intelligence - A Prescription for Operational Success by Claudia Imhoff (originally published October 23, 2007)

analysis into their operational processes and applications. This analysis is more specific and more common than traditional BI analysis, because it covers questions and decisions that are part of the organization's day-to-day business processes; for example, "What discount can I offer my clients based on their past ordering history?"

Operational BI focuses on optimizing an organization's performance in the present and not the past, and providing the organization with a near real-time view using dashboards, metrics, and reports. Traditional BI and operational BI are not the same, but they do complement each other. The traditional data warehouse can become a source of information to provide historical data to operational BI applications.

Components of an operational BI implementation is to include the following information:

- Historical data

  Self-service dashboards with links to traditional analysis and reports.

- Real-time data

  Information to include near real-time data that is also provided to the dashboards and reports.

- Business activity monitoring for complex event and data processing from both operational data sources and the data warehouse environment. Key events are distributed to those required using push technology, for example, email notification.

- Information on demand capabilities

  These capabilities are explained in *Enterprise Data Warehousing with DB2 9 for z/OS*, SG24-7637.

Operational BI is increasingly coming to resemble online transaction processing (OLTP). It is also becoming more critical that a BI application be reliable, continuously available, secure, and can process comfortably within a diverse workload. This shift toward operational BI and dynamic data warehousing requires predicable performance, not just acceptable performance, as the key to an enterprise's successful exploitation of business analytic tools.

These statements do not imply that performance is no longer important, because it is. Acceptable wait times for results continue to drop. Parallel query processing and advanced storage subsystem technologies, together with continually improving optimized database management systems, significantly improve the speed by which data can be accessed and evaluated. However, physical limits with traditional approaches call for more radical steps to satisfy tomorrow's business requirements.

In addition to predictable and improved performance expectations, the "people cost" to maintain a well-performing data warehouse is always increasing. Successfully maintaining a DW often requires a highly skilled and experienced DBA with access to a sophisticated set of tools to maintain the performance expectations of today's business analytics users. Traditional tuning methods, such as indexing and materialized query tables (MQT), are time-consuming and often deliver unpredictable results. Columnar tables, main memory databases supported by massive amounts of real memory, larger storage devices, and specialty engines all contribute to the complexities of managing today's data warehouse.

How can all of this be accomplished and still guarantee high qualities of service with unsurpassed security and scalability? This can be achieved by taking advantage of the one environment that has always been up for the challenge, the combination of System z, the z/OS operating system, and DB2 for z/OS. And now, this traditional System z solution is being extended to also provide more consistent query performance results.

From late 2005 to 2007, the Blink project in the IBM Research Division developed technologies to accelerate processing of long-running online analytical processing (OLAP) queries by orders of magnitude. Additional details can be found at:

`http://www.almaden.ibm.com/cs/projects/blink/`

The aim of this research project was to provide consistently low, interactive query response times, regardless of any specific query structures when accessing the underlying data. This is accomplished by exploiting the implementation of leading technology trends: hybrid row/column stores in main memory exploiting predicate evaluation on highly compressed data using multi-core and vector-optimized algorithms. The Blink project's technology, in combination with the IBM zEnterprise BladeCenter Extension (zBX), are the foundations for the development of the Smart Analytics Optimizer.

Data warehouses have been implemented on DB2 for System z servers largely because of the traditional strengths that the System z platform provides, such as unparalleled throughput, continuous availability, and unlimited scalability. The Smart Analytics Optimizer extends that by demonstrating the innovation that exists on System z.

The Smart Analytics Optimizer

The Smart Analytics Optimizer design point is to execute queries that are typically found in business intelligence and data warehousing applications with fast and predictable response times, thus offering and enhancing a comprehensive business intelligence solution on System z. Figure 20 shows the addition of the Smart Analytics Optimizer solution to a System z environment.

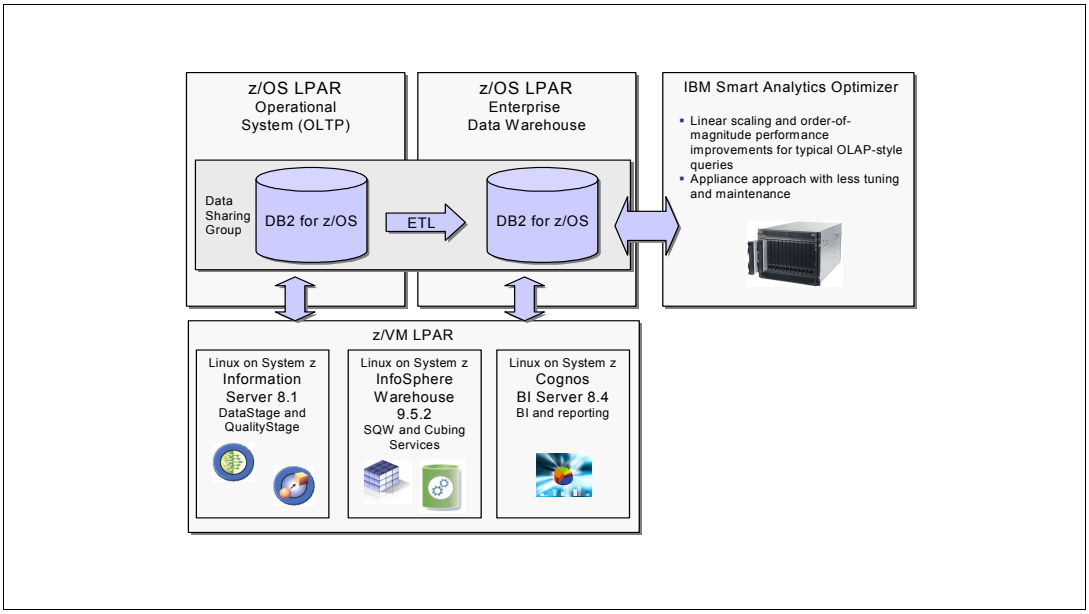

Figure 20 Adding the Smart Analytics Optimizer to System z

The Smart Analytics Optimizer is a combination of both hardware and software.

The software, Smart Analytics Optimizer for DB2 for z/OS, takes advantage of the attached BladeCenter to deliver the following benefits:

- Up to 10x performance gain on certain types of queries
- A more consistent query execution time for the varied BI workload

- Accelerated business insight for faster and better decision-making
- Reduced administrative time and cost associated with database tuning
- Higher qualities of service across fit-for-purpose, workload-optimized infrastructure
- Improved economics through better infrastructure management and resiliency

This solution is currently exclusive to zEnterprise 196 (z196) and z/OS. It consists of the following components:

- IBM BladeCenter attached to a z196 server.
- IBM software with installation features for simple initial deployment assisted by graphical tooling. This software is built from two major parts:
  - The unique code developed by IBM and named Smart Analytics Optimizer Application. The Smart Analytics Optimizer Application is a database-like application, based on Blink technology, containing IBM-patented compression algorithms.
  - A closed platform that runs on the BladeCenter and includes an operating system and the IBM General Parallel File System V3.x (GPFS™).
- DB2 for z/OS Stored Procedures running on a DBMS host system to which the Smart Analytics Optimizer solution is connected.
- Administration plug-in hosted by an IBM Data Studio Foundation client. The administration plug-in is named Smart Analytics Optimizer Studio. This Eclipse-based GUI runs on a workstation attached to System z and connected to DB2 for z/OS. It is used by database administrators for Mart definition and other administration activities.

The Smart Analytics Optimizer is an integrated hardware and software solution offering a centralized environment that extends System z legendary availability and security to heterogeneous business intelligence and data warehouse workloads; the first-of-a-kind workload optimized technology that is multiplatform.

Using operational controls on the Hardware Management Console (HMC) in combination with a BladeCenter, System z is capable of managing this heterogeneous BI infrastructure as a single entity with minimal risk and reduced costs. These benefits are provided without change to current applications, because DB2 for z/OS transparently takes advantage of the special purpose hardware and software for query execution by sending qualifying queries to the Smart Analytics Optimizer running on the zBX.

Although called the Smart Analytics Optimizer, this is not an “optimizer” in the way we have used the term optimizer in DB2. Rather, it is an attached query accelerator that relies on the traditional DB2 optimizer to make the decision as to whether to execute a query natively in DB2 or on the accelerator.

The hardware

IBM mainframe systems have been providing specialized hardware and dedicated computing capabilities for a long time. Not counting the machine instruction assists, recall the vector facility of the IBM 3090, back in the mid-1980s; the System Assist Processor, for I/O handling, which implemented the 370-XA architecture; the Coupling Facility and the Cryptographic Processors. And all the I/O cards, which are nothing less than specialized dedicated hardware with sophisticated software that offloads processing from the System z processor units (PUs).

The common theme with all of these specialized hardware capabilities is their seamless integration with the mainframe. They are configured, managed, and serviced the same way as the main CPC. Despite the fact that their processors might not be System z PUs and run

purpose-specific software, this software does not require any administration or tuning by customers; it is, in fact, invisible to them and is handled as Licensed Internal Code. These hardware features really are part of the mainframe, not an add-on.

As mentioned, Smart Analytics Optimizer is composed of both hardware and software. We look at the hardware in “Managing the zBX”and discuss software in “The software” on page 42.

To simplify the ordering and administration processes, Smart Analytics Optimizer is configurable in a variety of solution offerings based on raw data volumes that do not include DB2 indexes or DB2 materialized query tables (MQT). The zBX is connected to a zEnterprise 196 by OSA-Express3 10 gigabit Ethernet (GbE) connections for transporting data. The fiber connection uses TCP/IP and DRDA. An attached disk device is also required for persistent storage of the data marts on disk, before they get loaded into the memory of the zBX. Currently, this disk device is an IBM System Storage® DS5020.

Managing the zBX

A key feature of the zBX is its integration under the System z management umbrella. Thus, initial firmware installation, updates, and patches follow the pattern of System z. This also applies to the configuration and definitions.

Similar to channels and processors, the Support Element (SE) has an additional view for the blades in the zBX. This view shows icons for each of the zBX component’s objects including an overall status (power, operational, and so on).

The following functions and actions are managed and controlled from the System z HMC or SE:

- View firmware information for the BladeCenter and blades
- Retrieve firmware changes
- Change firmware level
- Backup/restore critical data
  - zBX configuration data is backed up as part of SE backup and restore on replacement of a blade.

For more details about the HMC and SE support for the zBX, refer to the Redbooks® publication *IBM zEnterprise System Technical Introduction*, SG24-7832.

The firmware for the zBX is managed, controlled, and delivered in the same way as for the System z server. It is packaged and tested with System z microcode and changes are supplied and applied with MCL bundle releases.

Benefits of the zBX firmware packaged with System z microcode are:

- It is tested together with System z driver code and MCL bundle releases.
- You can retrieve code using the same integrated process of System z (IBM RETAIN® or media).
- There is no need to use separate tools or connect to web sites to obtain code.
- You can utilize upcoming System z firmware features, such as Digitally Signed Firmware.
- The infrastructure incorporates System z concurrency controls where possible.
- zBX firmware updates are fully concurrent, using blades similar to Config Off/On controls.
- There is an audit trail of all code changes in the security log.
- There is automatic back out of changes to the previous working level on code apply failures.
- Accelerator firmware is provided.

The Smart Analytics Optimizer application, supporting operating system and management agent code, released as firmware with the rest of the code, is automatically downloaded (if necessary) from DB2 on the first connection.

The Smart Analytics Optimizer automatically assigns roles to the zBX blades. Each blade (node) can assume one of two role types: coordinator or worker.

Coordinator node Accepts and queues work. Distributes and coordinates work of the worker nodes and forwards results to requester (DB2).

Worker node Processes the work requests from the coordinator node. Note that there are spare nodes that can be used to back up the coordinator or worker node.

All available nodes in the zBX are used. If a worker node fails, one of the multiple coordinator nodes takes over for the failing worker node. If only a single coordinator node exists when the worker node fails, there is no failover and the zBX is then unavailable. After the original failing working node is replaced, the node previously acting as a worker node will be brought back up again as a coordinator node.

At any given time, a single query block can be processed on the zBX. The coordinator node can have other requests on the queue waiting to execute.

The numbers of coordinator and worker blades are defined by the specific configuration ordered.

The software

The IBM Smart Analytics Optimizer is optimized to mitigate database administration costs and speed up the performance of typical data warehouse workloads by orders of magnitude. To achieve these benefits, the workload must be sent to a zBX running the Smart Analytics Optimizer software.

The Smart Analytics Optimizer software is implemented as a logical extension of DB2 for z/OS, and thus is deeply integrated with the DB2 for z/OS Optimizer. It requires DB2 9 for z/OS or above (DB2 9 requires additional service), and z/OS 1.10 or above (z/OS 1.10 requires additional service). DB2 must run in *new function* mode, as illustrated in Figure 21.

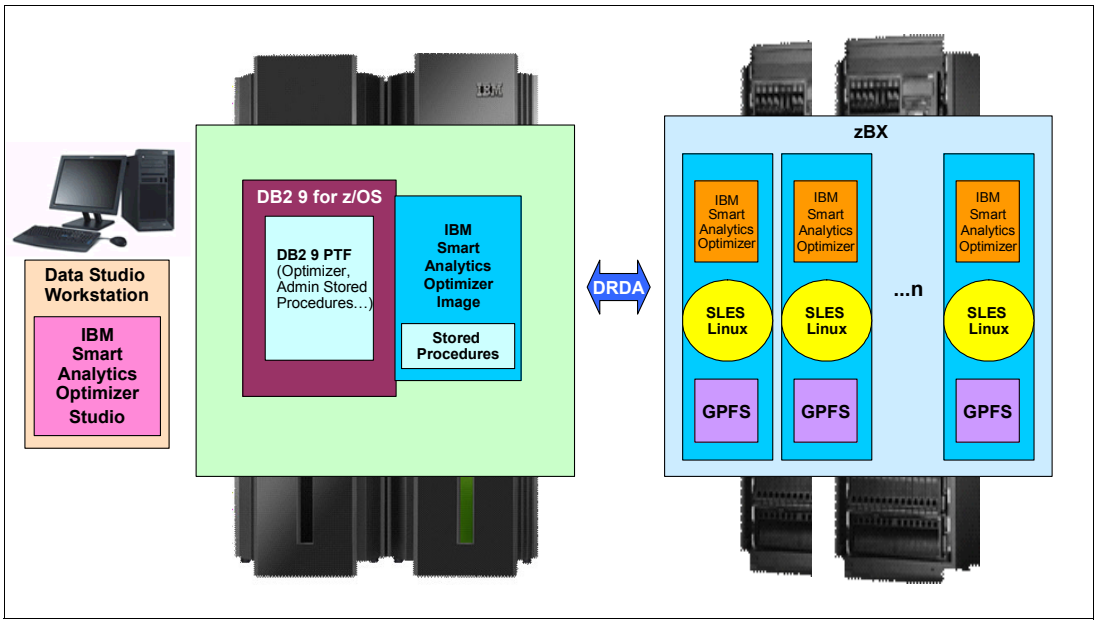

Figure 21 Smart Analytics Optimizer solution software components

The bundled software and the DB2 Stored Procedures comprise the software product Smart Analytics Optimizer for DB2 for z/OS. The Smart Analytics Optimizer software is shipped on a DVD. An SMP/E installation package provides integration with the DB2 for z/OS environment.

The Smart Analytics Optimizer software is installed on the zBX blades through a zEnterprise 196, using the product DVD. After it is installed, updates to the software are installed as PTFs and updated on the blades by calling a Stored Procedure.

The Smart Analytics Optimizer Studio software is installed from the DVD on a workstation that is attached to the System z server and connected to DB2 for z/OS.

DB2 for z/OS architecture

DB2 for z/OS consists of multiple resource managers: a log manager, a data manager, a buffer manager, and so on. These resource managers all exist within DB2 running under z/OS. They are also fully managed by DB2 for z/OS. The Smart Analytics Optimizer can be considered as an additional DB2 resource manager. However, in the case of Smart Analytics Optimizer, this particular resource manager resides on a zBX rather than natively in z/OS on a zEnterprise 196 server. This is shown in Figure 22.

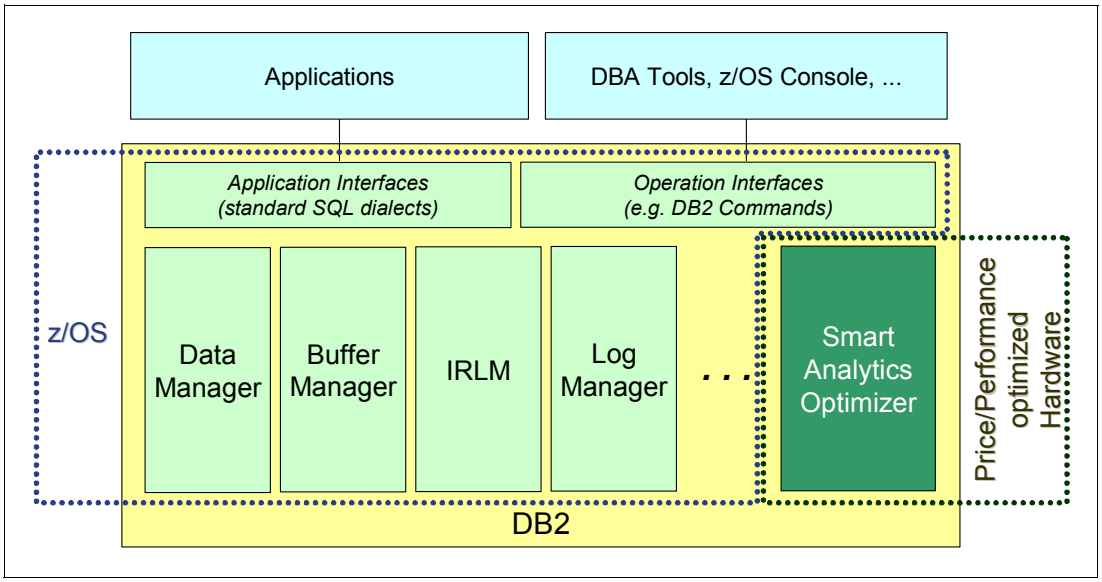

Figure 22 High-level architecture of the Smart Analytics Optimizer for System z

The Smart Analytics Optimizer hardware and software are managed completely by System z, z/OS, and DB2 for z/OS. There is no need for any outside interaction. In fact, there is no means provided to manage the Smart Analytics Optimizer without going through z/OS and DB2 for z/OS.

Workloads that qualify

Not all query types and data types are currently supported by the Smart Analytics Optimizer. The most noticeable processing exclusion from the Smart Analytics Optimizer in the first release is probably static SQL. Currently, only dynamic SQL is supported. Most BI analytics use a form of dynamic SQL, and dynamic SQL presents challenges when it comes to achieving acceptable performance and performance predictability.

There are restrictions in this release of Smart Analytics Optimizer. The following discussion lists various restrictions and other reasons that can cause DB2 to exclude a query type from processing in the Smart Analytics Optimizer.

Notably, the data referenced in the SQL statement's FROM clause must have already been loaded in the Smart Analytics Optimizer. Only tables that you have defined to be loaded in the Smart Analytics Optimizer can be used by a query. We will be discussing shortly how a table or set of tables are defined and loaded into the Smart Analytics Optimizer and how DB2's optimizer is aware of which tables are available for processing.

There are a few other restrictions in this release as well. No query can spread across multiple marts, and queries must include at least one fact table. As for join processing, only *inner joins* and *left outer joins* are allowed; full outer and right outer joins are prohibited.

There are certain DB2 data types that if used will exclude a query from taking advantage of the Smart Analytics Optimizer. They are binary types (including BINARY, VARBINARY, and BLOB), decimal floating point, XML, a row identifier (ROWID), CLOBs, and DBCLOBs. As you will see, one method for getting a query to qualify to execute on the Smart Analytics Optimizer when one or more of these data types are used by a column definition is to exclude that column from the table when the mart stored in the Smart Analytics Optimizer is created.

In addition to these data types, certain DB2 functions are not allowed. Although most built-in SQL functions are supported, there are functions that, if specified, will prevent an SQL statement from being routed to the Smart Analytics Optimizer. Functions that will currently cause an SQL query to *not* be selected for routing to the Smart Analytics Optimizer include:

- User-defined functions
- Mathematical functions such as SIN, COS, TAN, and EXP
- Character-wise string functions
- Advanced string functions including LOCATE, LEFT, OVERLAY, and POSITION (just to name a few)
- Advanced OLAP functions like RANK, DENSE, ROW NUMBER

Note that routing a query to the Smart Analytics Optimizer is the responsibility of the DB2 optimizer. The DB2 optimizer makes the decision that a query qualifies for the Smart Analytics Optimizer using a mechanism called the accelerated query table (AQT).

The primary reasons that a query can be precluded from executing on the Smart Analytics Optimizer include:

- The query is a statically bound statement.
- The CURRENT REFRESH AGE special register is set to 0 (zero).
- The query contains SQL syntax not supported by the Smart Analytics Optimizer.
- The accelerator or mart has been disabled.
- The query references a table or table column that is not loaded in a defined accelerator mart, perhaps due to a table that is not defined in the mart, a column being referenced that was excluded from the marts definition, or the column data type being unsupported.
- The query does not reference a fact table.
- The DB2 optimizer makes the decision that the query can be executed natively on DB2 and achieve equal or better performance. Even though the Smart Analytics Optimizer is available, DB2 cost-based estimations still apply.

There are two types of queries, usually associated with business intelligence, that form the “sweet spot” for the Smarts Analytics Optimizer: *star* and *snowflake schemas*, and *complex ad hoc* queries. In addition, as previously mentioned, only inner joins and left outer joins where the fact table is on the preserved side are allowed, and full outer and right outer joins are prohibited. Analysis of query workloads determined that by far, the dominant analytical queries are inner joins and left outer joins. There were few full joins contained in the supplied workloads.

Note that this is the first release of Smart Analytics Optimize, and various items discussed might be supported at a future time.

Marts

It is important to understand that not all data of a data warehouse is stored in main memory of the Smart Analytics Optimizer. Only those tables that have been identified by the DBA as targeted by queries that qualify to execute on the Smart Analytics Optimizer. Eligible queries for the Smart Analytics Optimizer are mostly long-running OLAP queries. Short-running OLTP queries will seldom if ever qualify. All queries must reference the fact table loaded across the worker nodes in the zBX.

Those tables being accessed by long-running OLAP queries will be sent as a snapshot to the Smart Analytics Optimizer. What is identified and stored in the Smart Analytics Optimizer is a set of tables which are related to each other. One example is a star or snowflake schema. In data warehouse environments, you typically have a large fact table and two or more smaller dimension tables. One set of those tables is referred to as a *mart*, and a mart is technically what is going to be sent to the Smart Analytics Optimizer. Figure 23 shows a mart being stored in the Smart Analytics Optimizer.

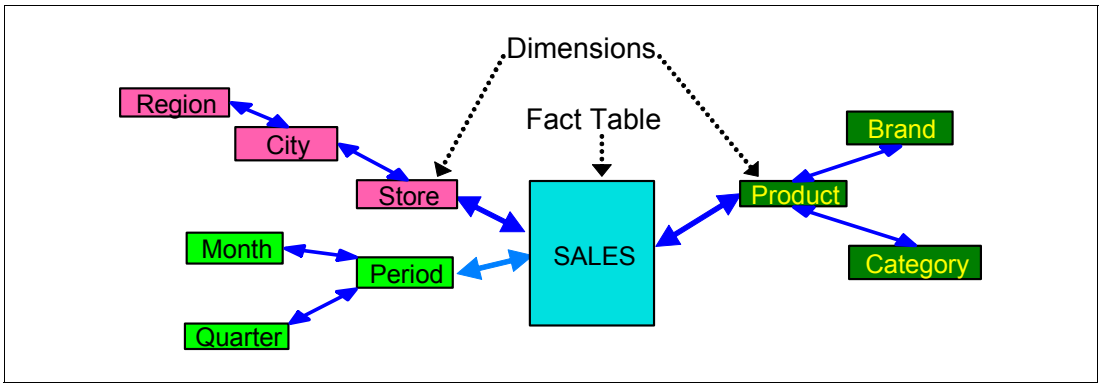

Figure 23 A mart example for the Smart Analytics Optimizer

A mart is a collection of logically related tables centered around one fact table. If there are multiple fact tables, there will be multiple marts. What is actually stored in the Smart Analytics Optimizer is a mart. The set of tables belonging to a mart can be identified using a graphical user interface (GUI). That graphical interface available through the Smart Analytics Optimizer Studio, an extension of the no-charge Data Studio product, is used to define those relationships.

After identifying the relationships between the tables, either by automatically detecting referential constraints or by manually defining the relationships between those tables, a mart can be loaded into the Smart Analytics Optimizer and the definition of the marts stored in the Smart Analytics Optimizer own catalog tables. Unloading the data from the DB2 table used to load the data into the marts is performed by using the UNLOAD utility and the utilization of UNIX System Services pipes to send the data to the Smart Analytics Optimizer.

Smart Analytics Optimizer Studio can also be used to include the most often referenced columns, excluding columns that are seldom or never used. This method can be employed to minimize the size of the marts stored in the zBX. The smaller the marts, the more marts can be stored in the zBX.

Usually, only a subset of all data stored in a data warehouse is related to long-running OLAP queries. Therefore, the total size of all data warehouse data does not necessarily indicate how many tables can be stored in the Smart Analytics Optimizer and how many queries of a given installation can be accelerated. It is best to define and send as a mart to the Smart Analytics Optimizer only those portions of data that can greatly benefit from a dramatic increase of performance.

Note that data stored in a data warehouse within DB2 for z/OS is usually compressed. Because the Smart Analytics Optimizer uses another compression algorithm from DB2 for z/OS, you will observe more data being sent through UNIX System Services pipes during the load process to the Smart Analytics Optimizer than is currently stored on disk in your DB2 for z/OS subsystem. Additionally, there can be a separate compression ratio for the same data within the Smart Analytics Optimizer compared to DB2 for z/OS.

Unloading data for the mart for the Smart Analytics Optimizer can exploit parallelism, with the degree of parallelism being determined at Smart Analytics Optimizer installation. The degree is also influenced by the number of partitions. As the data is unloaded, it is written to a named pipe which is read by another stored procedure, then sent through TCP/IP to the Smart Analytics Optimizer load mart code. Compression of data stored at the Smart Analytics Optimizer occurs in the zBX. Make sure that enough temporary space is available to store the data sets that are unloaded to the Smart Analytics Optimizer.

When the process begins to unload the data to the Smart Analytics Optimizer, portions of the data is stored across all available blades, thus allowing for massive parallel processing on the blades. The Smart Analytics Optimizer can store as many marts as capacity allows it to, and within a given query only data residing in one mart is allowed to be accessed by a query that is routed to the Smart Analytics Optimizer.

This means that no SQL statements utilizing the Smart Analytics Optimizer are allowed that cross a mart's boundary. Figure 24 shows a definition of two marts using the same dimension tables but other fact tables. In this example, if a query can access both fact tables, the query is not eligible to be routed to the Smart Analytics Optimizer.

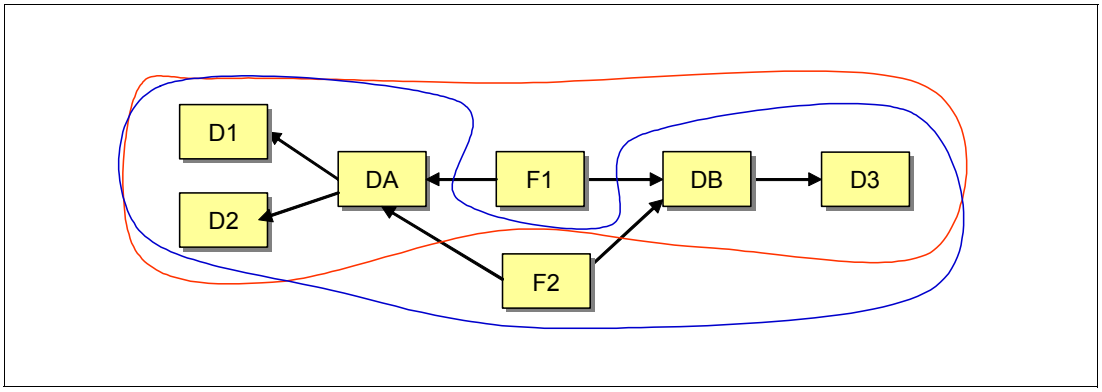

Figure 24 Multiple mart definitions

Whenever a mart is defined, the information about that specific mart is stored in the DB2 for z/OS catalog, enabling the optimizer to take these definitions into account when calculating the most efficient access path for a given query. The decisions made by the optimizer for routing queries accessing data that is also stored in the Smart Analytics Optimizer are based on a cost estimate, as for traditional DB2 for z/OS workloads. From an Smart Analytics Optimizer perspective, a mart can be one or more Accelerated Query Table.

An AQT has characteristics that are similar to an MQT and carries many of the same restrictions. One significant difference between an MQT and an AQT is that the AQT definition does not carry all of the local predicates with it as an MQT does. For this reason, there is a much higher likelihood that more queries will qualify against it.

Note that the Smart Analytics Optimizer is designed to accelerate long-running OLAP queries. It is not designed to accelerate short-running OLTP queries. It is important to understand that a snapshot of data is stored and queried in the Smart Analytics Optimizer, as is common for DW environments. The original data will still remain in DB2 for z/OS. Accessing data either in the Smart Analytics Optimizer or from disk is controlled by the special register CURRENT REFRESH AGE.

Because DW and BI data in most cases is aggregated before providing the final results of a query, then depending on your specific situation, it is likely that the result of a query is not going to be affected in a way that the result changes significantly if non-snapshot data is queried. This depends on the type of query and type of data you are going to use in your SQL

statements. If small deviations in the result set can be neglected, resulting from incremental changes to your data warehouse data before the mart is refreshed with the most current data, it is useful to have all long-running OLAP SQL statements enabled to be routed to the Smart Analytics Optimizer.

Because the Smart Analytics Optimizer stores a snapshot of parts of your DW data, it can be important for specific queries to access the data as it is stored within DB2 for z/OS because most current data can be mission-critical for specific queries. However, as with MQTs, it is also possible to query the data as it is stored within DB2 for z/OS even if the query is eligible to be routed to the Smart Analytics Optimizer. Whether or not the optimizer will consider an existing mart definition within the Smart Analytics Optimizer for query optimization is determined by the value of the special register CURRENT REFRESH AGE (as with MQTs):

- ANY enables the Smart Analytics Optimizer usage for the optimizer.
- 0 disables the Smart Analytics Optimizer usage for the optimizer.

Example 1 shows the syntax to enable and also to disable Smart Analytics Optimizer usage for a query succeeding the statement shown within the same thread.

Example 1 Enabling or disabling Smart Analytics Optimizer usage for a given query

```
Enable Smart Analytics Optimizer usage:     SET CURRENT REFRESH AGE = ANY;
Disable Smart Analytics Optimizer usage:    SET CURRENT REFRESH AGE = 0;
```

Development and deployment process

Before queries can be accelerated on the Smart Analytics Optimizer, you need to follow a definition and deployment process (Figure 25). The key activity in this process is defining marts, and the key tool to use is IBM Data Studio. Before starting this process you need to understand not only which marts are required and what is the query behavior, but also which queries can benefit from running on the Smart Analytics Optimizer and thus need a mart to be defined. Performing an assessment, as explained in “Smart Analytics Optimizer assessment” on page 49, can be a useful starting point for gaining this insight.

The definition step identifies the tables and relations that comprise marts. In the deployment step the marts are made known to DB2; for example, storing mart metadata in the DB2 and Smart Analytics Optimizer catalog. Smart Analytics Optimizer Studio guides you through the process of defining and deploying marts, and in invoking other administrative tasks. Smart Analytics Optimizer Stored Procedures implement and execute various administrative operations such as mart deployment and load and update. They serve as the primary administrative interface to Smart Analytics Optimizer from the outside world, including Smart Analytics Optimizer Studio.

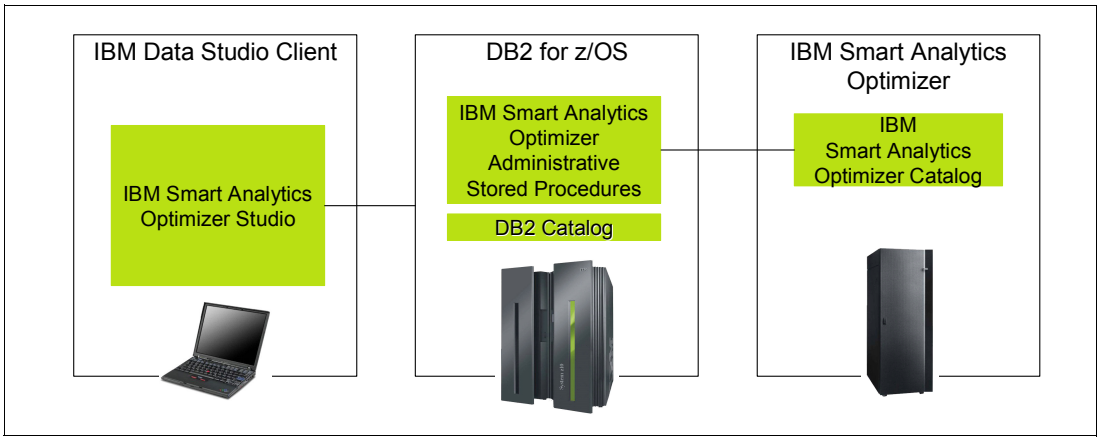

Figure 25 Smart Analytics Optimizer definition and deployment process

Tooling

IBM Smart Analytics Optimizer Studio (Figure 26) is an administration plug-in hosted by the IBM Data Studio client, part of a non-fee edition of the Data Studio line of products. This Eclipse-based GUI is used by database administrators for data mart definition and other administration activities.

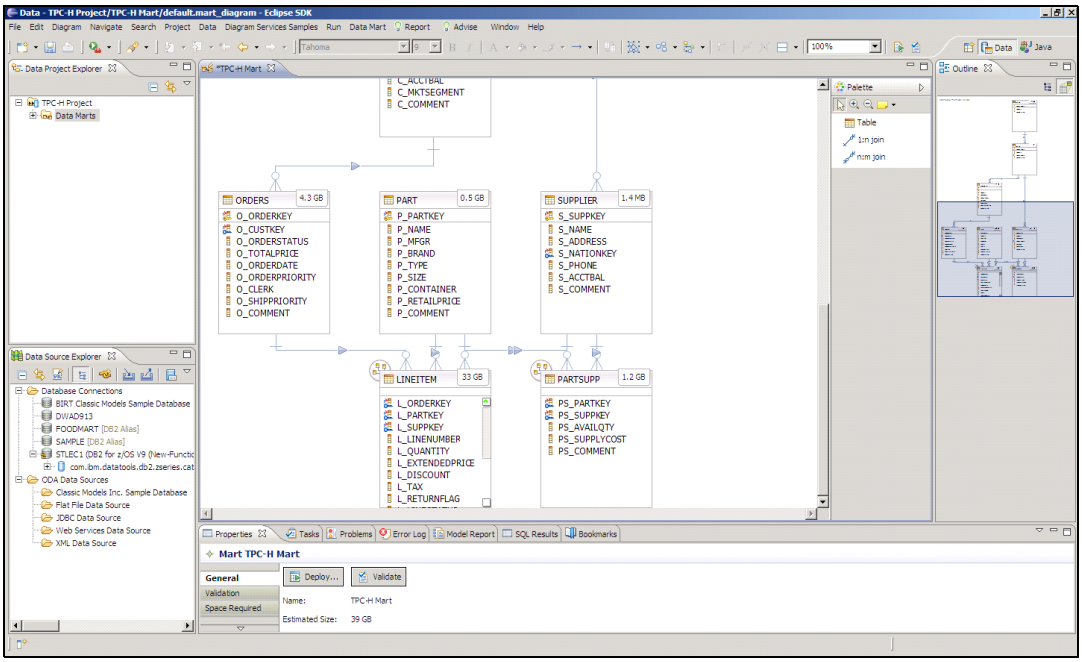

Figure 26 IBM Smart Analytics Optimizer plug-in

Smart Analytics Optimizer assessment

The System z Center of Excellence at the development lab in Boeblingen, Germany has devised an assessment tool that can help you determine if any of your queries qualify and can take advantage of the Smart Analytics Optimizer.

Because the Smart Analytics Optimizer processes only dynamic queries in its first iteration, the DB2's dynamic statement cache makes an excellent starting place for determining which portion of a client's query workload might potentially qualify. To take advantage of tracking queries using the statement cache, the dynamic statement cache must first be enabled. This task is completed by setting the DSNZPARM keyword CACHEDYN on the DSN6SPRM macro or the CACHE DYNAMIC SQL field on the DSNTIP8 installation panel to YES. Although the feature is designed to capture dynamically prepared SQL for potential reuse, enabling the dynamic statement cache results in all dynamic SQL being captured.

If you are using the statement cache to only capture SQL for the assessment, consider clearing the cache prior to running the warehouse queries for the assessment. The dynamic statement cache can be cleared by running the DB2 utility RUNSTATS using the REPORT NO UPDATE NONE control cards, issuing an SQL DROP or ALTER against an object used by an SQL statement in the statement cache, recycling DB2, or revoking the SQL authorization for the authid responsible for getting the SQL statement into the statement cache.

The next step required when preparing to perform a Smart Analytics Optimizer assessment is turning on the correct DB2 traces to ensure that all of the necessary information about the dynamic SQL statements being captured in the statement cache is available. You accomplish this by turning on a trace that will gather the IFCID 316 and 317 records. In a non-data sharing environment, use the following command:

```
-START TRACE(MON) CLASS(1) IFCID(316,317,318) DEST(SMF)
```

In a data sharing environment, use the following command:

```
-START TRACE(MON) CLASS(1) IFCID(316,317,318) DEST(SMF) SCOPE(GROUP)
```

You can use the **-DIS TRACE(*)** command to verify that the correct traces are active.

The IFCID 316 trace record collects the first 60 bytes of the SQL statement, and identifying information and statistics. The IFCID 317 record is used in addition to IFCID 316 to obtain the full SQL statement text. IFCID 318 acts as a switch for the IFCID 316 and 317 records. Use IFCID 318 to toggle the 316 and 317 records active and inactive. Finally, overhead from running these traces is minimal with no more than 4% anticipated when the IFCID 318 record is active. IFCID 318 is to be inactive when not actively gathering information with the statement cache.

Stopping and starting IFCID 318 can be used to start a new trace interval for the 316 and 317 trace records. Ensure that the traces for IFCID 316, 317, and 318 are started prior to collecting any data warehouse query workload, and that they remain active until after the EXPLAIN STMTCACHE statement completes successfully. If these traces are not active throughout the entire process of capturing the dynamic SQL and dumping the statement cache, then all the information necessary to complete the assessment will not be present.

The number of dynamic SQL statements captured to complete an assessment can have an impact on the amount of SMF records collected. Discuss this with the z/OS systems programmers or the person responsible for collected SMF data to avoid collection problems caused by additional SMF overhead.

Note that dumping the dynamic statement cache to the EXPLAIN statement cache table has no affect on the contents of the statement cache. Some action, like running RUNSTATS or altering an object used by a dynamic SQL statement in the cache, must be performed to clear the cache, as previously mentioned.

Ensure that the dynamic SQL being collected is for a time interval that best reflects a typical query workload. Dynamic SQL collection needs to last long enough that a true representative DW/BI dynamic SQL sample is collected. After you are satisfied an adequate SQL sample

has been gathered, you move on to performing a “dump” of the dynamic statement cache. Keep in mind, however, that the dynamic statement cache is of fixed size. When the cache is full, recording does not stop but instead wraps and records over the first records saved in the cache. If a sample is gathered over an extended period of time, the number of dynamic SQL statements gathered can be less than anticipated if the dynamic statement cache is too small for the interval.

After a query workload sample has been collected, you dump and EXPLAIN your dynamic SQL. First, however, verify that ALL EXPLAIN tables have been created. You can use Optimization Service Center (OSC), Optim™ Query Tuner, and sample job DSNTIJOS in DB2's SDSNSAMP PDS to create the appropriate required EXPLAIN tables. DSNTIJOS is shipped with the Smart Analytics Optimizer assessment tool.

The following EXPLAIN tables are used by the assessment process:

- PLAN_TABLE
- DSN_STATEMNT_TABLE
- DSN_FUNCTION_TABLE
- DSN_PREDICAT_TABLE
- DSN_STRUCT_TABLE
- DSN_QUERY_TABLE
- DSN_STATEMENT_CACHE_TABLE

Also used is the CLOB column for statement text in XML format from DB2 9:

- DSN_DETCOST_TABLE

With all of the proper EXPLAIN tables in place, you can EXPLAIN all of the dynamic SQL statements that have been captured in the dynamic statement cache. Make one more check of the EXPLAIN to ensure that the SQLID (or SCHEMA name) is correct. After you verify this, then using either SPUFI, DSNTEP2, or any other application that allows the execution of dynamic SQL, execute the SQL statement EXPLAIN STMTCACHE ALL.

This SQL statement populates the DSN_STATEMENT_CACHE_TABLE (referred to here as the cache table) using the current contents of the dynamic statement cache. Run a SELECT SQL statement against the cache table to validate that the contents reflect the dynamic queries you were attempting to capture. This is also an opportunity to eliminate, using a qualified SQL DELETE statement, any SQL statements from the cache table that are not representative of the warehouse query workload you are attempting to capture.

With the EXPLAIN's DSN_STATEMENT_CACHE_TABLE table populated from the statement cache, run the supplied REXX exec against the cache table. The REXX exec will perform an EXPLAIN against each of the SQL statements in the cache table. After the exec completes successfully, these EXPLAIN tables will have the information needed to complete the assessment. Using a set of jobs supplied as part of the assessment, the EXPLAIN tables and selected DB2 Catalog tables are unloaded, tersed, and transmitted to IBM for final analysis.

From the final analysis, a paper is prepared and delivered to the client summarizing the assessment results. This results report details the top 50 queries, which queries qualify and how much of the queries qualify, along with a list of the tables used by the queries analyzed by the assessment.

You can arrange for an assessment by contacting your IBM representative.

Summary

This guide summarizes the System z value proposition for information management in general and focused on the additional value of the IBM Smart Analytics Optimizer. System z, combined with the Smart Analytics Optimizer, provides a complete foundation for full-function information management architectures with high performance and high reliability.

Other resources for more information

The following IBM Redbooks publications contain further information about the topics discussed in this guide.

On demand capabilities and data warehousing:

- *Enterprise Data Warehousing with DB2 9 for z/OS*, SG24-7637

HMC and SE support for the zBX:

- *IBM zEnterprise System Technical Introduction*, SG24-7832

Parallel Sysplex technology and benefits:

- *Workload Management for DB2 Data Warehouse*, REDP-3927

The team who wrote this guide

This guide was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO).

Alex Louwe Kooijmans is a Project Leader with the International Technical Support Organization (ITSO) in Poughkeepsie, NY. He specializes in SOA technology and solutions on System z. He also specializes in application modernization and transformation on z/OS. Previously he worked as a Client IT Architect in the Financial Services sector with IBM in The Netherlands, advising financial services companies on IT issues such as software and hardware strategy and on demand services. Alex has also worked at the Technical Marketing

Competence Center for zSeries® and Linux in Boeblingen, Germany, providing support to customers getting started with Java and WebSphere® on System z. From 1997 to 2000, he completed a previous assignment with the ITSO, managing various IBM Redbooks projects and delivering workshops around the world in the areas of WebSphere, Java, and e-business technology on System z. Prior to 1997, Alex held a variety of positions in application design and development and in product support and project management, mostly in relation to the IBM mainframe.

Willie Favero is an IBM Senior Certified IT Software Specialist with the IBM Silicon Valley Lab Data Warehouse on the System z Swat Team. He has more than 30 years of experience working with databases, and more than 25 years working with DB2. Willie presents at major international conferences and user groups, publishes articles and white papers, and has a technical blog on the Internet.

Fabricio Pimentel is a DB2 for z/OS and tools specialist in IBM Latin America at the System z Center in São Paulo Brazil. His areas of expertise include data management, performance, and availability. Fabricio holds a Bachelor's degree in Computer Science from Pontifical Universidade Catolica de São Paulo.

Thanks to the following people for their contributions to this project:

Bill White
ITSO, Poughkeepsie Center

Patric Becker
Information Management DWH on System z Center of Excellence, IBM Boeblingen Lab, Germany

Oliver Draese
Architect - IBM Smart Analytics Optimizer, IBM Boeblingen Lab, Germany

Guy Lohman
Manager of Disruptive Information Management Architectures, Advanced Database Solutions (K55), IBM Almaden Research Center

Alan Meyer
Marketing Manager Data Warehousing System z, IBM Software Group

Caryn Meyers
Marketing Manager Business Intelligence on System z

Bill Schray
Advanced Technical Support (ATS) for DB2

Mathias Zapke
User Experience Lead for IBM Smart Analytics Optimizer, IBM Boeblingen Lab, Germany

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author - all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in

length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document, REDP-4606-00, was created or updated on April 8, 2011.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at
http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

BladeCenter®
Cognos®
DataStage®
DB2®
Distributed Relational Database Architecture™
DRDA®
ESCON®
FICON®
GPFS™
HiperSockets™
IBM®
IMS™
Information Agenda™
Informix®
InfoSphere™
Optim™
Parallel Sysplex®
QualityStage™
RACF®
Redbooks®
Redguide™
Redbooks (logo) ®
RETAIN®
Sysplex Timer®
System Storage®
System x®
System z10®
System z®
Tivoli®
WebSphere®
z/OS®
z/VM®
z10™
zSeries®

The following terms are trademarks of other companies:

Java, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.