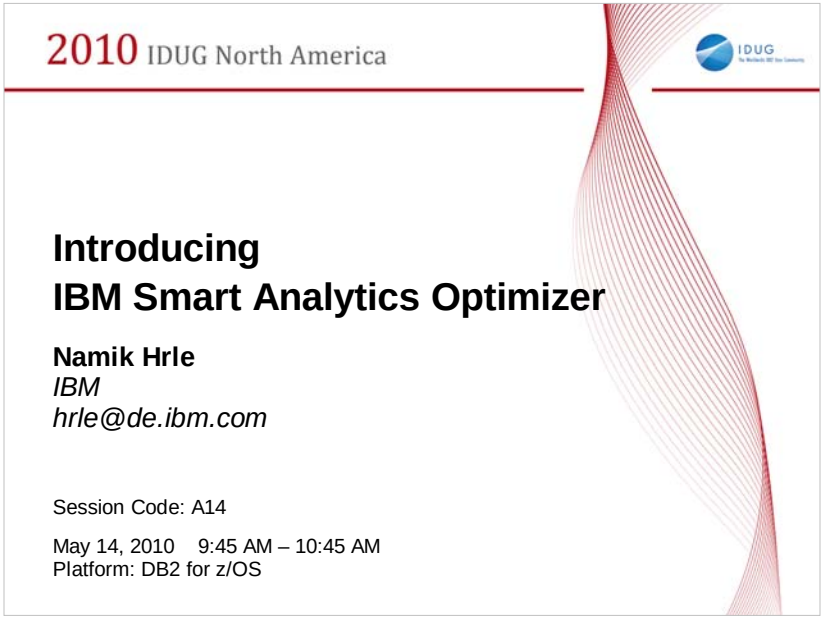

IBM Smart Analytics Optimizer (ISAO) is an appliance that seamlessly connects to DB2 for z/OS and dramatically accelerates selected data warehouse queries. Your Business Intelligence and Data Warehouse application do not need to be modified - they continue to access DB2 and are completely unaware of the existence of the accelerator. The DB2 Optimizer has been enhanced to decide whether to execute the queries in traditional way (in DB2) or to redirect them to the ISAO. The ISAO augments DB2 for z/OS and makes it the prime choice for BI and DW workloads. This presentation will give you a technology preview of the exciting new product, from the rationale for using it to the basic understanding of how it functions.

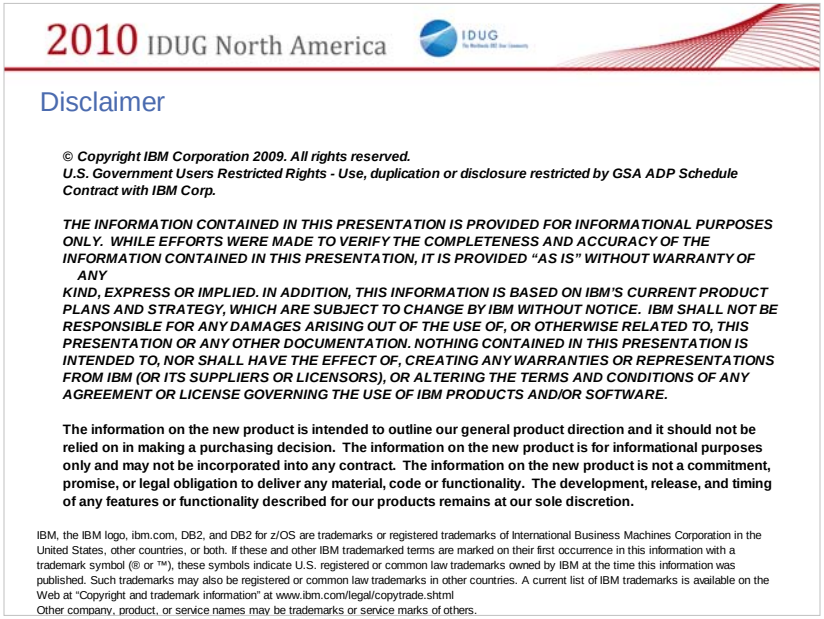

This product is in development and although we started the Beta program no assumption should be made about its general availability date, packaging, pricing etc. Any function can be withdrawn from the product at any point in time.

This is a technology preview only.

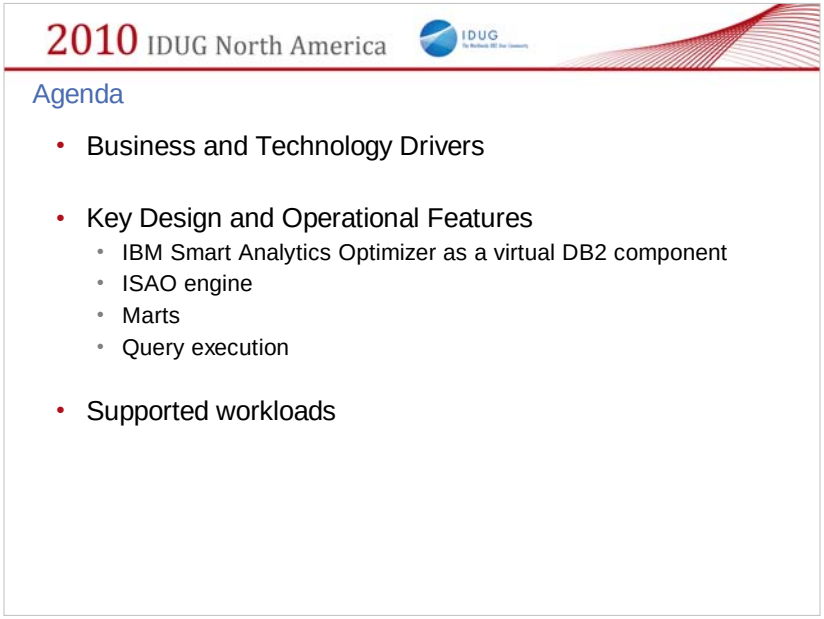

Five bullet points:

(1) Present business and technology drivers for introducing ISAO

(2) Understand ISAO place in DB2 landscape

(3) Explain ISAO query redirect capabilities

(4) Understand ISAO content maintenance

(5) Provide guidance for rough assessment whether ISAO applies to your specific environment

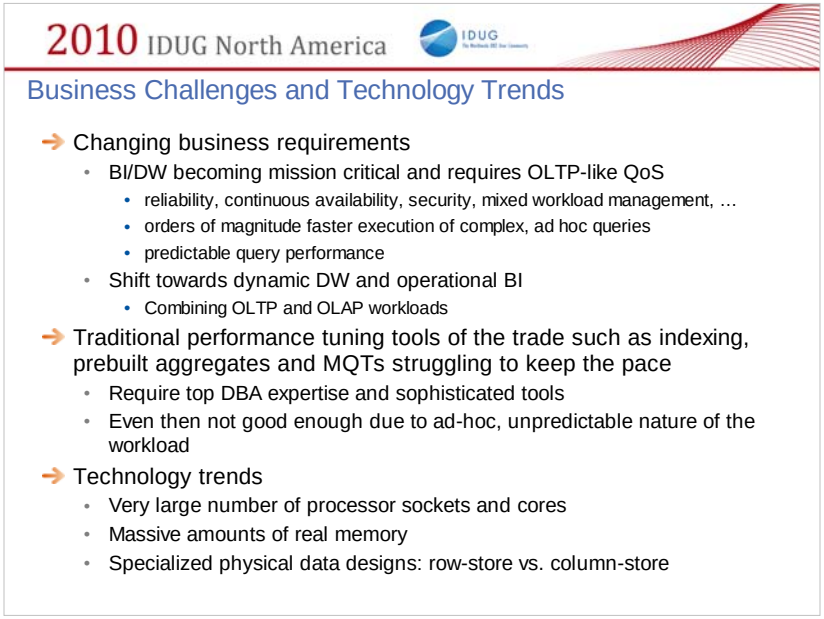

The trend of Business Intelligence and Data Warehousing becoming mission critical is being accelerated. These workloads are closing the gap that traditionally existed and becoming more like OLTP in regards to the requirements of being continuously available, secure and reliable.

Much faster query execution is expected in order to support new usage patterns of BI/DW applications. General purpose databases such as DB2 can provide very good performance but it comes with a considerable cost of continuous monitoring and tuning apart from requiring top expertize and sofisticated tooling, meaning significantly increasing TCO of the solution.

Traditional techniques of appropriate indexing, pre-aggregation etc. cannot really keep the pace with ever changing workload and unpredictable, ad-hoc queries.

All this is happening in conditions of very interesting technology trends massive scale out, availability of very large number of processors and cores measured in hundreds and thousands, and massive amount of real memory measured in Tbs. This enables a fresh new look on how to design software to take advantage of the aboundance of resources, so e.g. Column store databases are becoming more popular for warehousing.

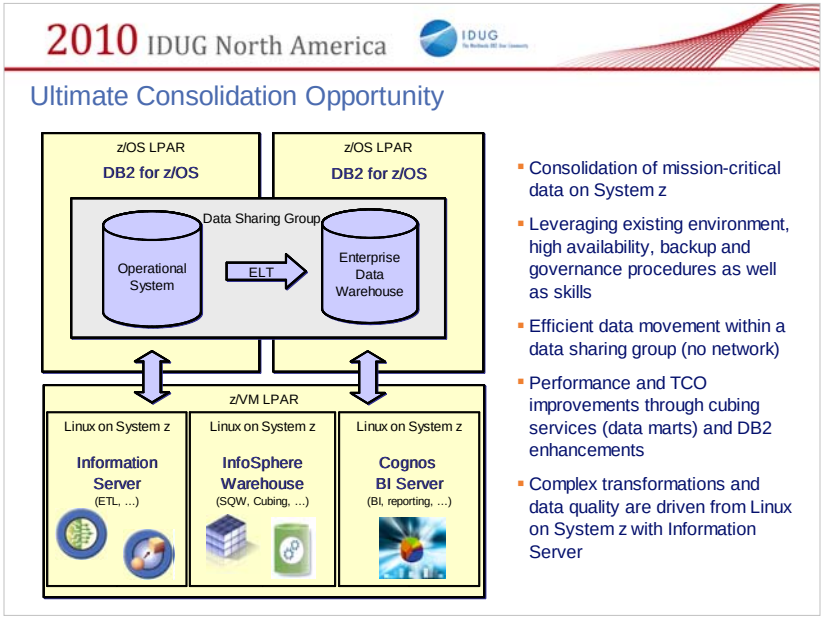

Apart from reliability, continuous availability and security the capability to handle mixed workloads is considered as one of the main platform differentiators for hosting modern data warehousing.

All this sounds like a great match for System z value proposition with its proven reliability, continuous availability and security features and industry leading workload management capabilities. On top of that, a very large amount of OLTP generated data is already managed by System z which opens possibility for exploiting synergic effects of proximity of operational and warehousing data.

Recognizing these business drivers and System z opportunity in the process of consolidating mission critical data irrespective whether operational or warehouse, we have brought the entire BI/DW portfolio to System z: Information Server with components like DataStage and QualityStage, Infoosphere warehouse (SQW and Cubing Services) and of course Cognos as the industry leading BI product. Combined with proven System z capability to host various workloads with the highest level of isolation especially when deployed in Parallel Sysplex and Data Sharing environment, this enables creating the ultimate consolidation opportunity, leverage existing skills, tooling, availability, governance procedures etc. The data can be moved between the data warehousing without leaving the box, without network and the latency of ETL can be significantly reduced.

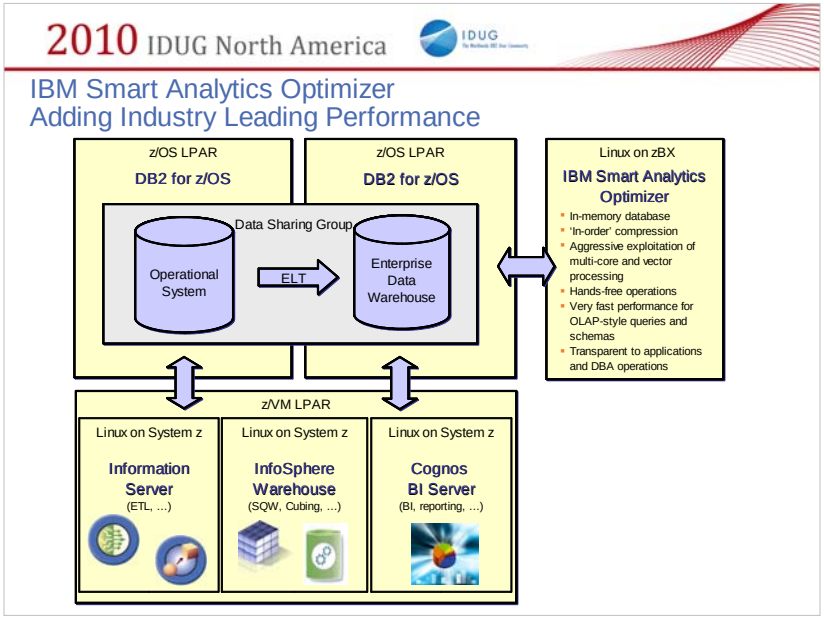

In order to prove the industry leading performance we have developed IBM Smart Analytics Optimizer (ISAO), which takes advantage of the technology trends and delivers orders of magnitude faster performance with virtually no tuning needs. At the same time the applications that access data do not need to change, they continue to connect to DB2 and transparently use the new technology.

Adding ISAO to DB2 for z/OS creates the ultimate DBMS that can handle both the OLTP and DW workloads and excell in both of them.

The rest of this presentation is devoted to the box you see on the right hand side of the chart, the ISAO.

Agenda

- Business and Technology Drivers
- Key Design and Operational Features
  - IBM Smart Analytics Optimizer as a virtual DB2 component
  - ISAO engine
  - Marts
  - Query execution
- Supported workloads

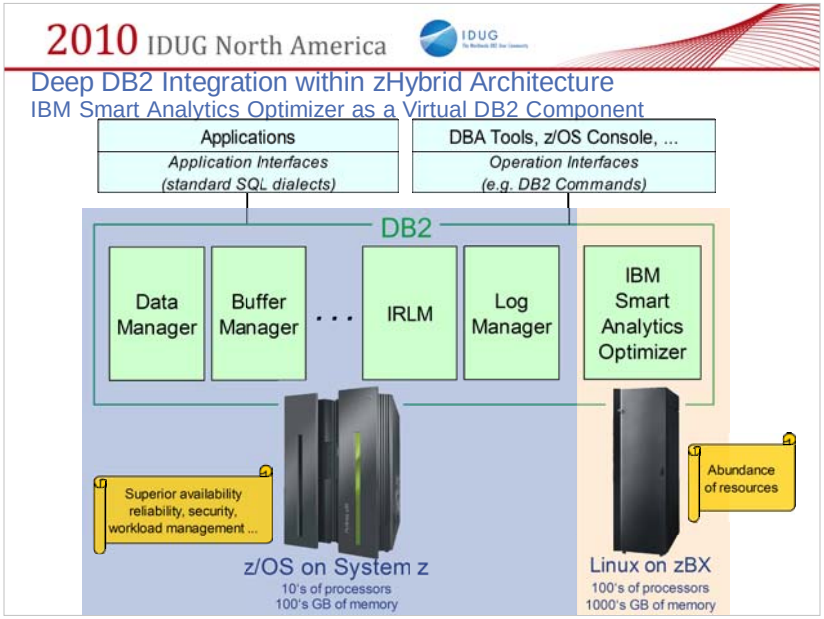

While all other DB2 components are implemented on a common code base and operating system services (in DB2 case that's z/OS), the SAO is an out-board component. By the nature of its workload, it requires massive parallelization and huge memory, that in terms of price/performance cannot be realistically implemented on System z..Namely, ISAO is a single workload server, whilc System z main strength is its legendary capability to handle many and mixed wor´kloads. So, it does not make sense to take System z and make out of it a single purpose system.

So, we decided to host that new DB2 component on a platform that offers abundance of resources and at the same time does not need all the System z quality of service characteristics.. ISAO will be implemented, packaged and installed as an appliance, with no interactions with the operational staff, no tuning, no direct log-on access in everyday operations

So, addition of the SAO creates a hybrid platform which is very much on the strategic path of future enterprise systems. The key take away from here is that the non-z elements of the hybrid are transparent to the outside world.

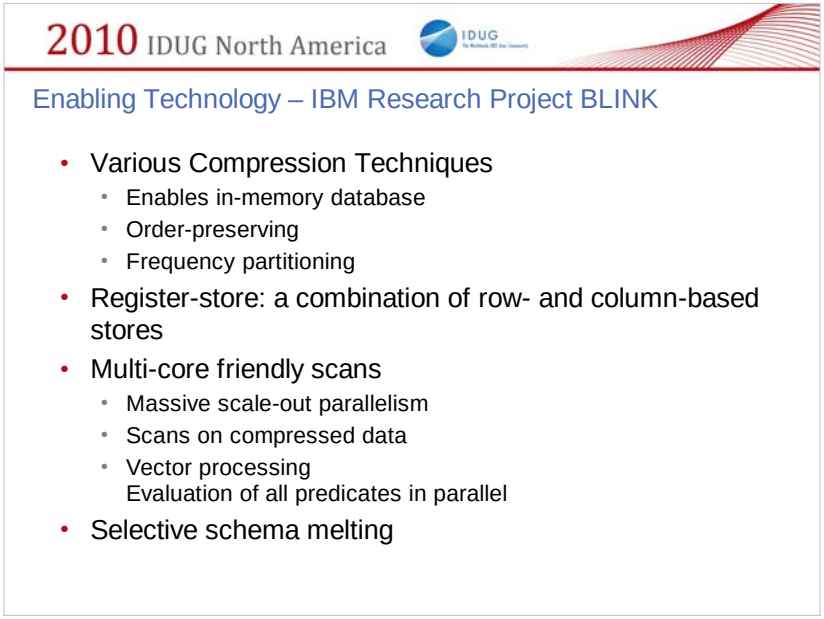

The ISAO internal engine is based on the IBM Research work. It 's a combination of various technologies: extreme compression, massive parallelism, register-store etc.

Extreme compression enables ISAO as an in-memory database. So, our scan algorithms assume that the entire set of tables included in the query resides in memory. The memory is distributed across theoretically unlimited number of blades – the memory availability dictates the number of required blades. The extreme compression includes numerous innovative compression techniques (many patents pending), such as frequency partitioning and fixed-length encoding, but the most notable characteristic is that they are order-preserving. The corollary of this is that the scans are done on compressed data – there is no need to decompress them for predicate evaluation. Decompression is done only for aggregation which is typically a small part in the entire processing flow.

Now that we have the data in memory and that we do not need to spend time in decompressing it before the scans and predicate evaluation, we can deploy the next technology – massive scale-up and –out parallelism. The algorithms that we deployed are very much multi-core and L2 cache aware, highly optimized for massive parallelism and processor, large registers and large L2 cache characteristics. There is very little data exchange across the nodes and we are able to achieve linear scaling. That means that adding nodes for growing database size maintains linear scaling and constant response time irrespective of the database size. The algorithms are heavily vectorized using SIMD which enables simultaneous application of predicate conjucts in only one compare operation. We always scan the entire table set whereas concurrent users can share a scan in L2 cache. Consistent response time.

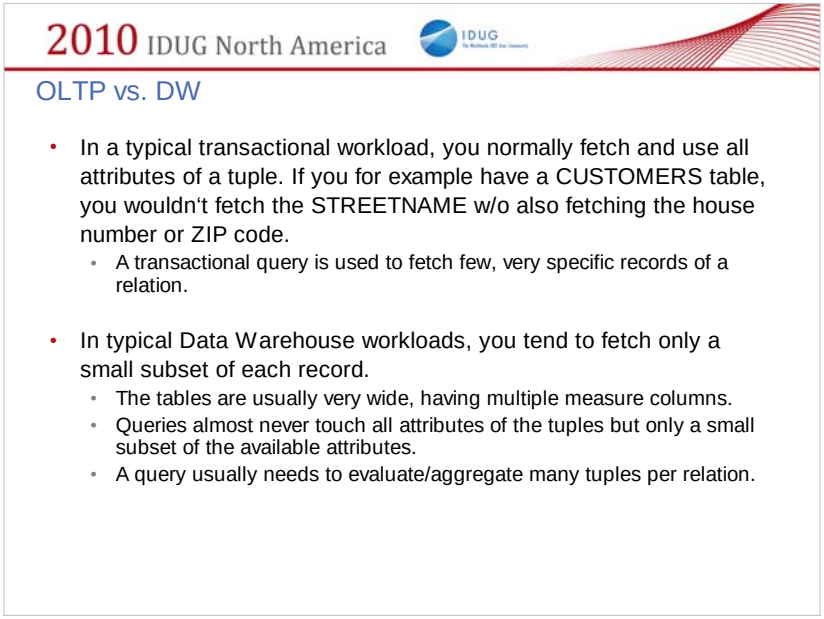

Typically, the type of queries that can be found in OLTP systems are different from those in the DW applications. This leads to different optimization techniques that are described on the following pages.

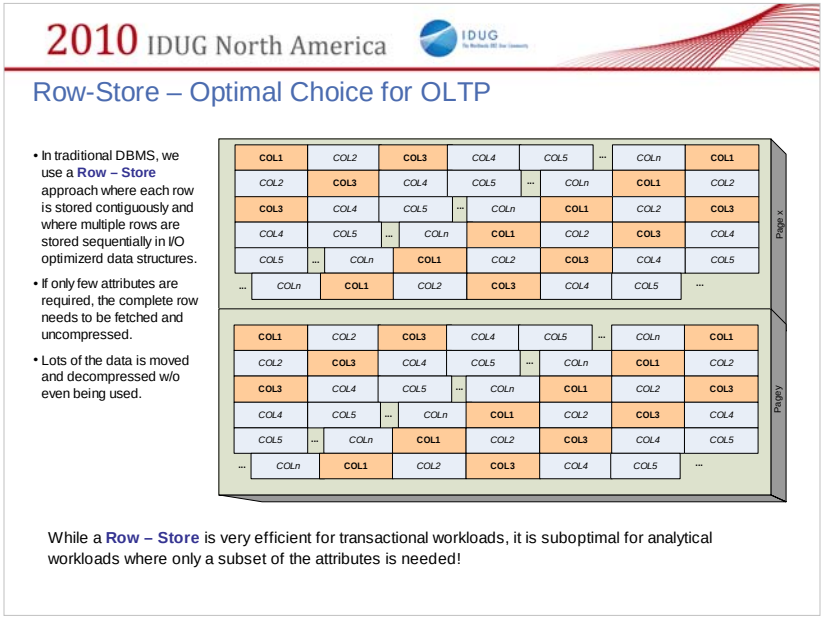

Row-stores are used in most general purpose DBMSs such as DB2, Informix, Oracle, SQLServer etc. The values of all the columns of row 1 is followed by the values of all columns of row 2, etc. This store model is very efficient for frequent updates to the database as well as the queries that can be optimized by using various indexing techniques.

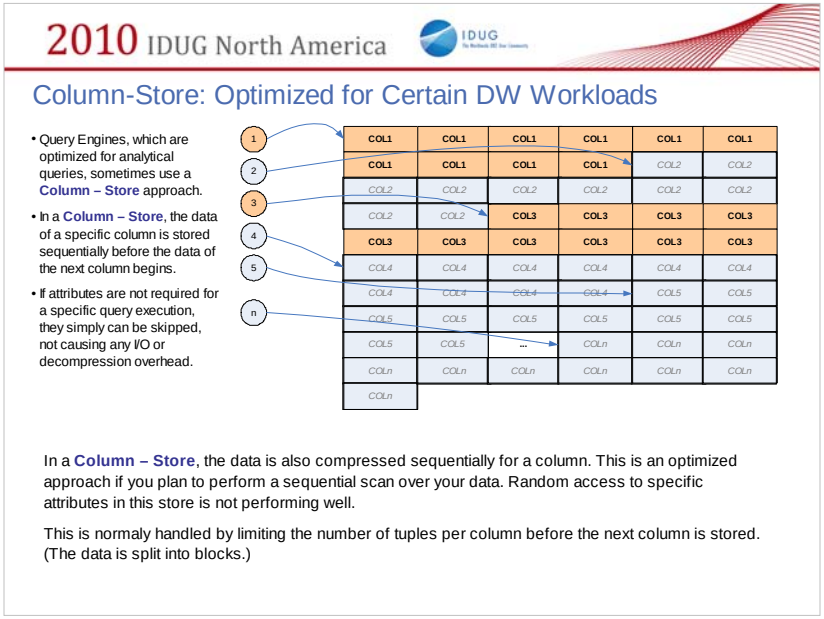

Column stores are popular in special purpose data warehouse oriented DBMSs. The values of column 1 of all the rows is followed by the values of C2 of all the rows etc. This store model is very efficient if only limited number of columns needs to be returned and very large number of data needs to be scanned.

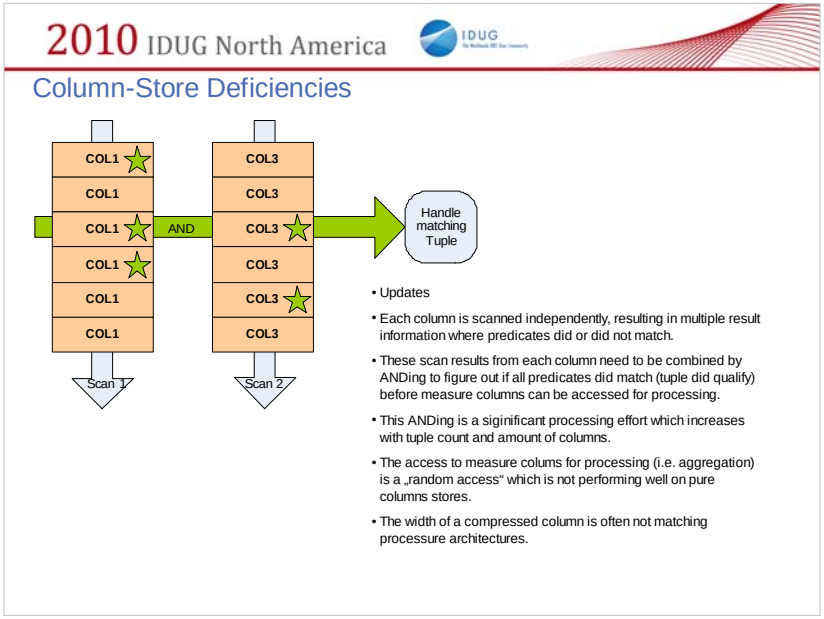

Of course, there are operations that are not efficient with column stores. They are listed on this chart.

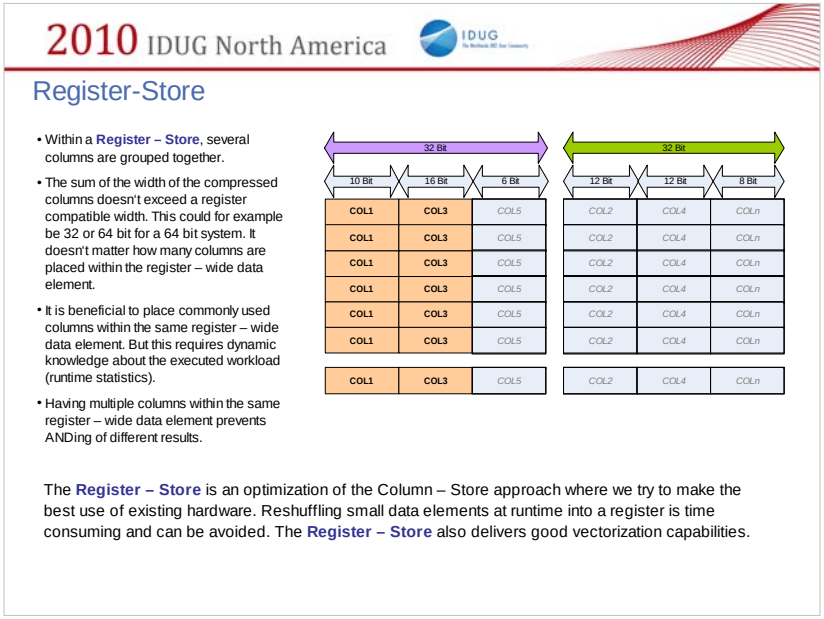

ISAO implements a hybrid between row- and column-store. We call it register-store because the processor register is defining element in deciding which columns will be grouped together. The groups of columns can be then very efficiently evaluated with as few processor cycles as possible.

ISAO uses so called 'bin packing' mechanism to decide which columns belong to which group.

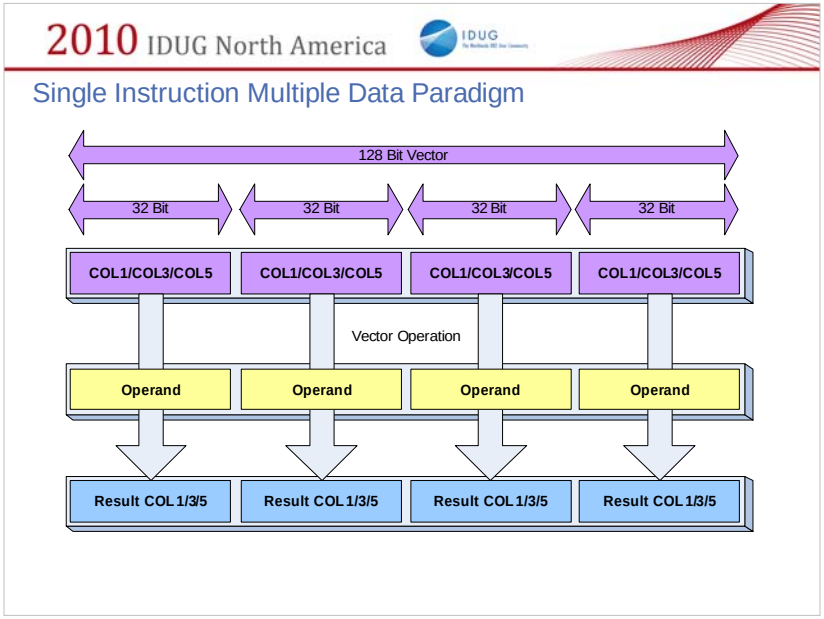

Additionally to column grouping and processor efficiency that is derived from that, ISAO also uses vector processing (Single Instruction Multiple Data operators) to perform simultaneous predicate evaluation on as many rows as possible.

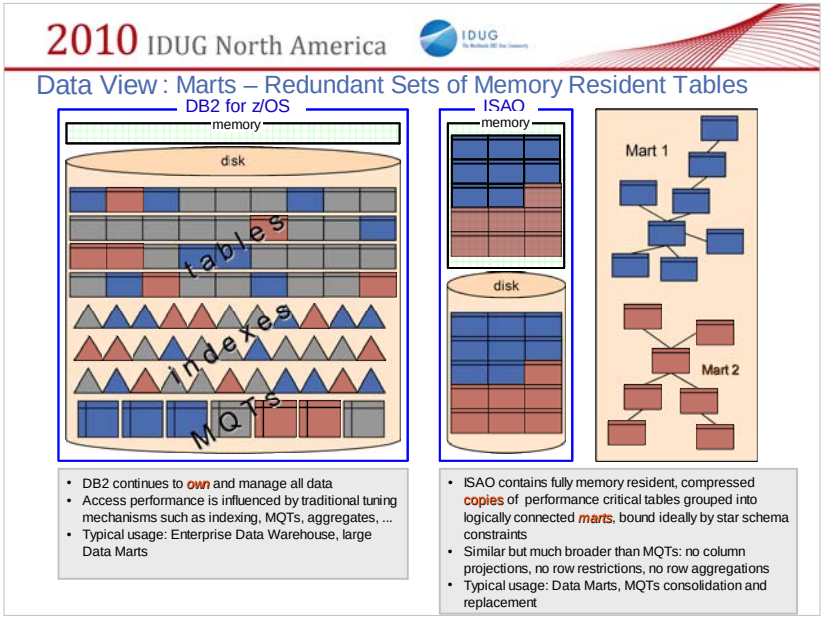

A mart is the key element of the ISAO data model. It consist of a number of logically connected tables, typically a fact table surrounded by dimension tables. It helps to think of marts as materialized query tables 'on steroids'. Namely, unlike MQTS, the marts do not include restriction and aggregatino operators and only optionally the can include the projection operators. Therefore, a mart can cover much broader set of queries than an MQT.

ISAO includes copies of the marts in highly compressed and scan optimized format.

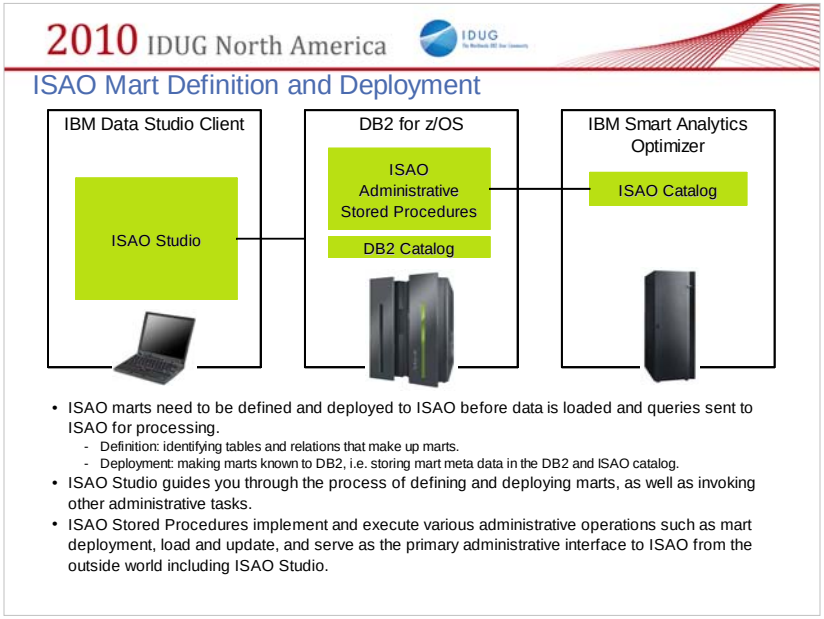

IBM Smart Analytics Optimizer Studio, an Administration plug-in hosted by an IBM Data Studio client (a part of no-fee edition of the Data Studio line of products). This Eclipse-based GUI is used by database administrators for data mart definition and other administration activities.

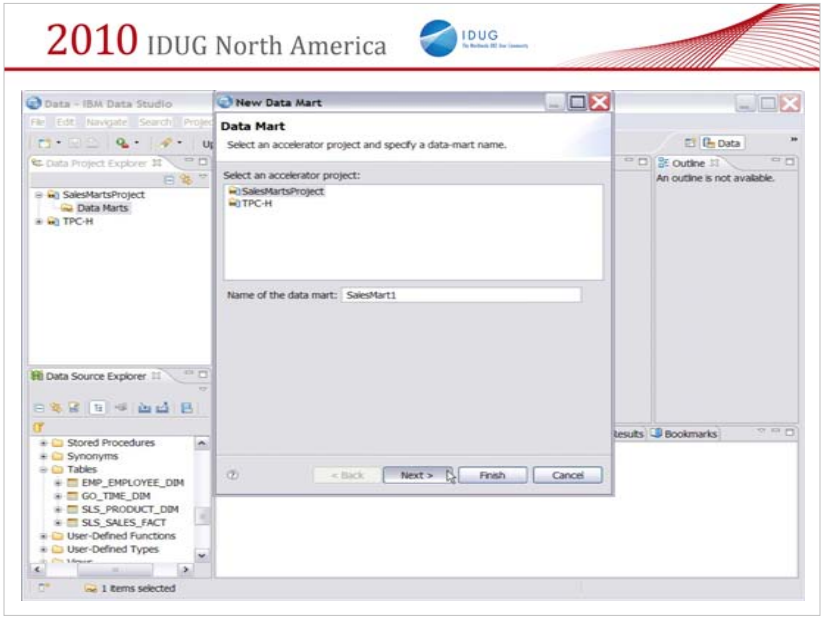

The following charts demonstrate a process of creating a mart using the ISAO Studio.

We start by creating a new project or by selecting an existing project. The project is a context in which we will define marts.

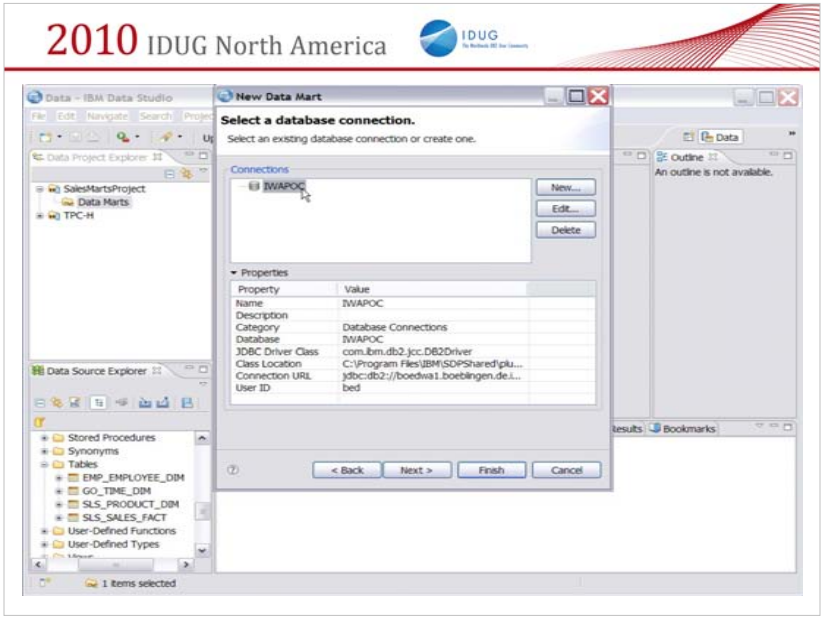

On this chart you have an option to select for which DB2 the mart will be created. At the same time you supply the basic connection parameters,

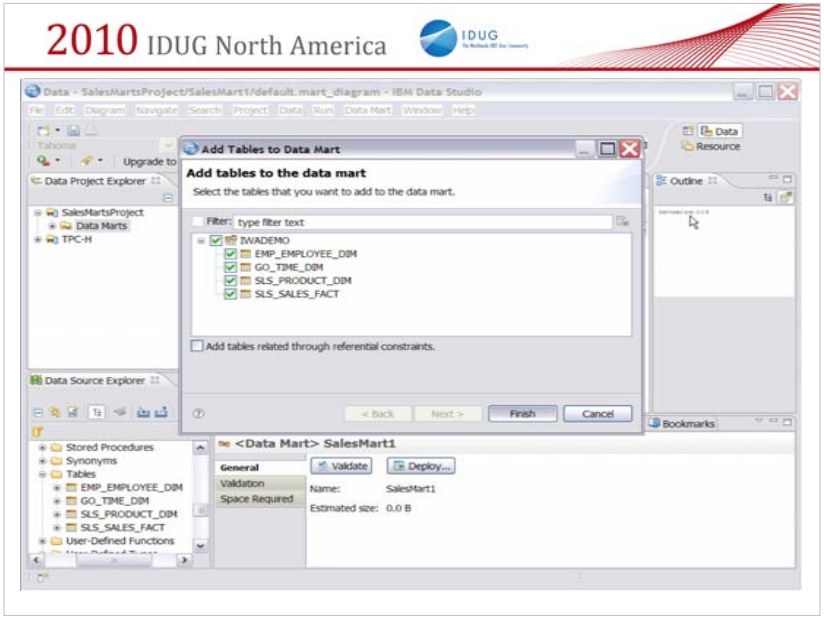

Now, you need to specify which tables will make up a mart. The tool gives a list of all available tables and you need to select those that are logically connected, i.e. those referenced by the queries that need to be accelerated.

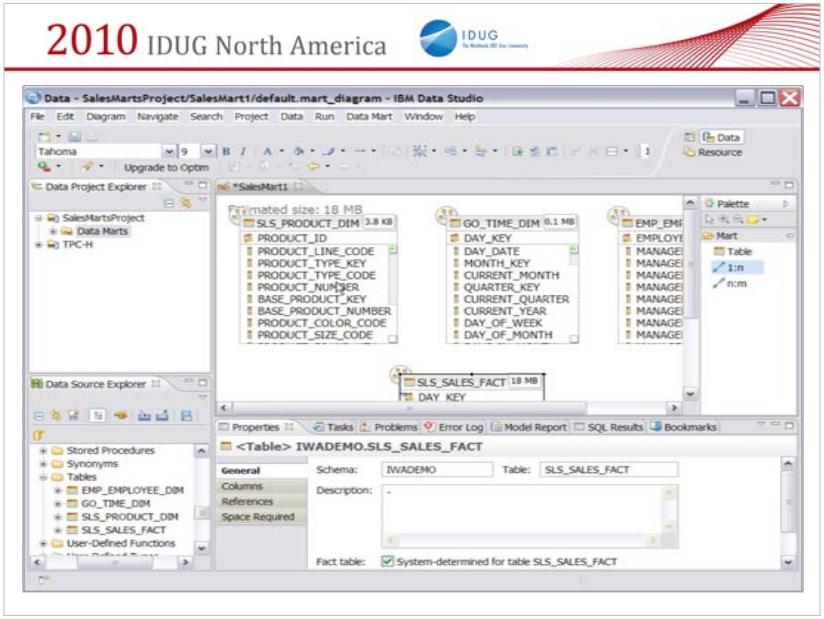

The tool presents the selected tables in the central pane and provides size estimates. The size is the amount of memory each of the tables is estimated to need in the ISAO.

At the same time a tool for defining relations among the tables is presented in the upper, right pane.

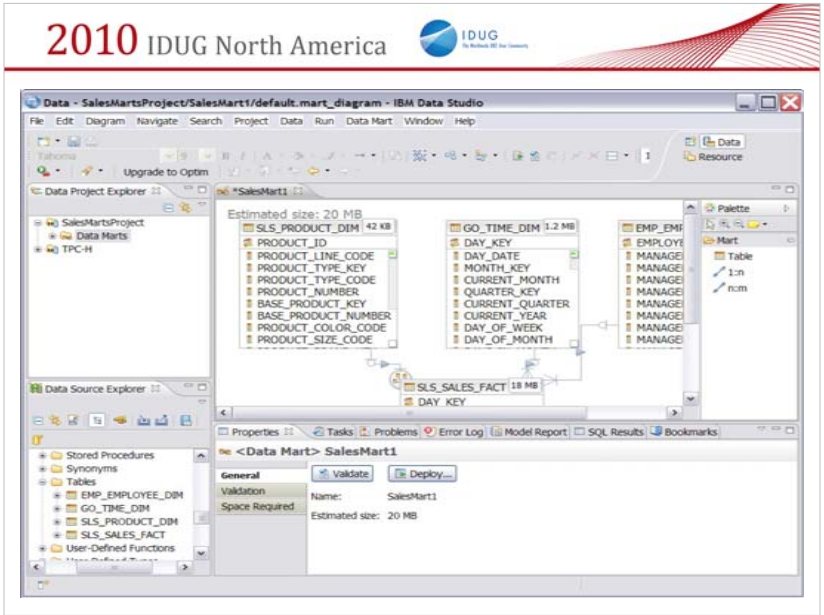

If the referential integrity rules are explicitly specified in DB2, the tool connects the corresponding tables automatically. Otherwise the relations need to be entered manually.

This completes the mart definition. The next step is the mart deployment (registering the mart in DB2 and ISAO), followed by the mart load.

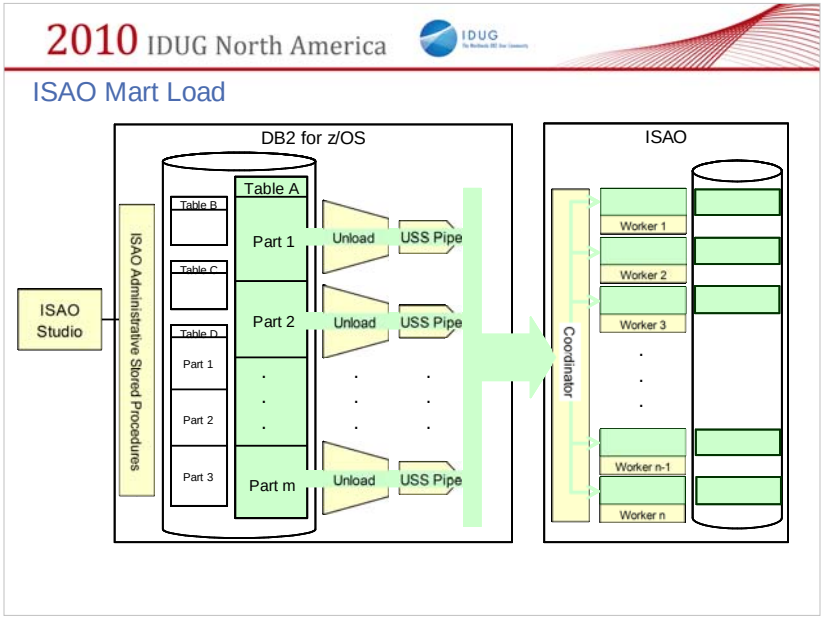

Marts need to be imported into ISAO. The initial load is done in a highly optimized way, using multiple parallel processes that pipe data through z/OS and TCP/IP (the data is never materialized on the System z side). Once at the ISAO side, the data is distributed to the workers by the coordinator. The workers compress the data, create scan optimized data structures in their memory and backup the memory content to the local ISAO disk. The backups are used in the case of a blade failure for fast repopulating of the memory content into the failover blade.

2010 IDUG North America

ISAO Mart Update

- Typical DW update operations:
  - LOAD RESUME and REPLACE
  - ADD and ROTATE PARTITION
  - SQL INSERT, UPDATE, DELETE
  - Delete complete partition or table
  - TRUNCATE TABLE
- ISAO will over time phase-in support for all the typical operations in this order
  1. Full table reload
  2. Updated partition reload
  3. Individual row change
- The marts update is initiated and controlled through ISAO Studio
- Queries off-loaded to ISAO before the marts are refreshed can return different result set as compared to not being off-loaded
  - In case this is not acceptable use SET CURRENT REFRESH AGE = 0

IDUG study on DW update frequency

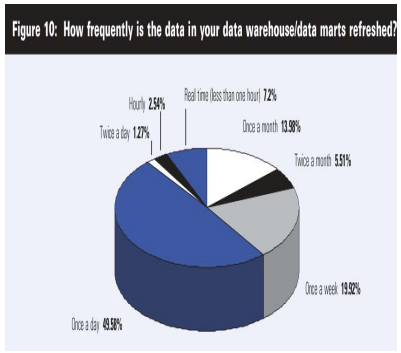

In 90% of cases DW is updated once a day or even less frequently

The data in the mart does not have to be refreshed all the time. There are operations that require only incremental updates.

The first release of ISAO supports a full refresh and more granular updates will be added subsequently.

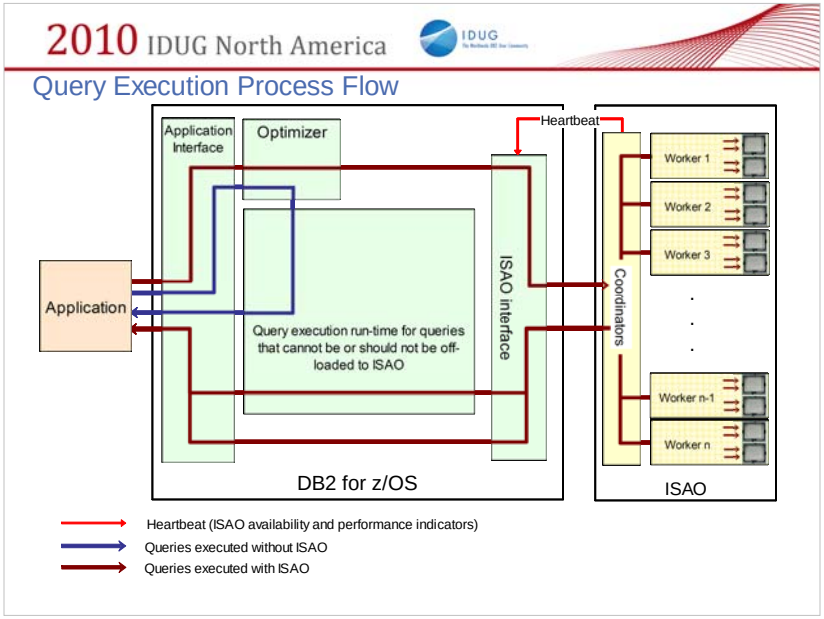

DB2 Optimizer has been enhanced to re-route eligible queries (cost based optimization) to the ISAO.

The queries that are better off being executed in DB2 will continue to be executed in DB2. That creates a hybrid environment that can handle both the OLTP and OLAP like queries.

In some cases the results from ISAO will be returned to the application without further processing by DB2, but in some cases the query blocks executed in ISAO and DB2 are combined within DB2. In any case, the application is not aware of ISAO presence.

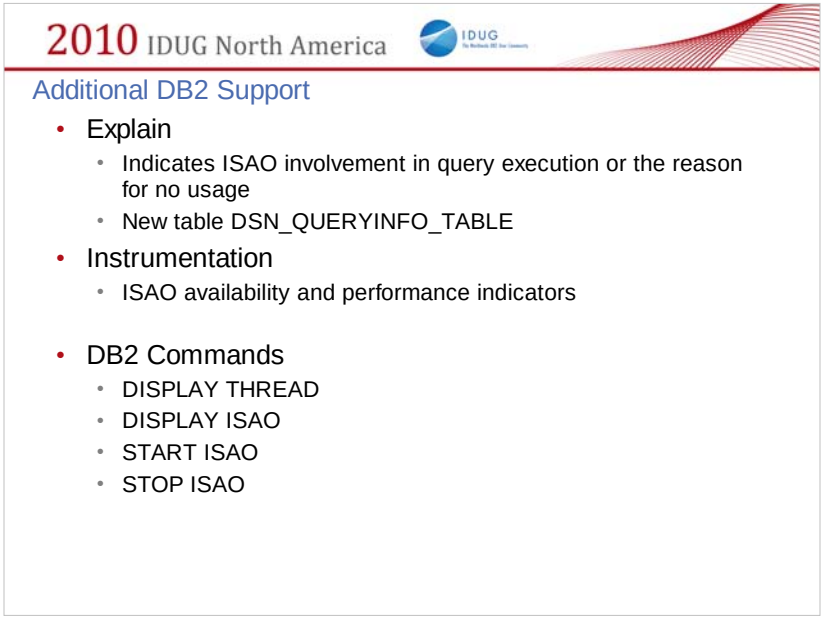

DB2 host also includes few externals relevant to ISAO. E.g. The Explain function has been enhanced to include descriptions for the blocks executed by the ISAO.

Possible ISAO Packaging Options

System z10

BladeCenters

ISAO Offerings	XS .5TB	S 1TB	M 2TB	L 3TB	XL 4TB
Enterprise Rack	1	1	1	2	2
BC-H Chassis, incl. double infrastructure (AMM, PDU, 10Gb Enet for data, 1Gb Enet for service, 4Gb F/C switches)	1	1	2	3	4
Blades	7	14	28	42	56
DS5020 with # of SATA disks (750GB)	16	16	16	32	32
Storage Expansion Units	0	0	0	1	1

DS5020

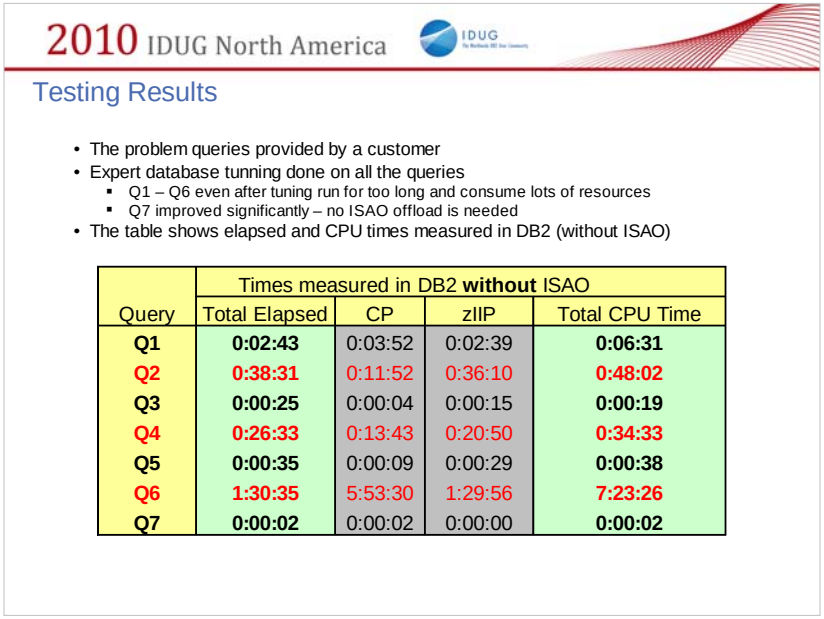

These 7 queries have been picked by the customer as the top queries in terms of elapsed time and CPU utilization. One of them (Q7) could have been tuned for efficient DB2 execution without ISAO.

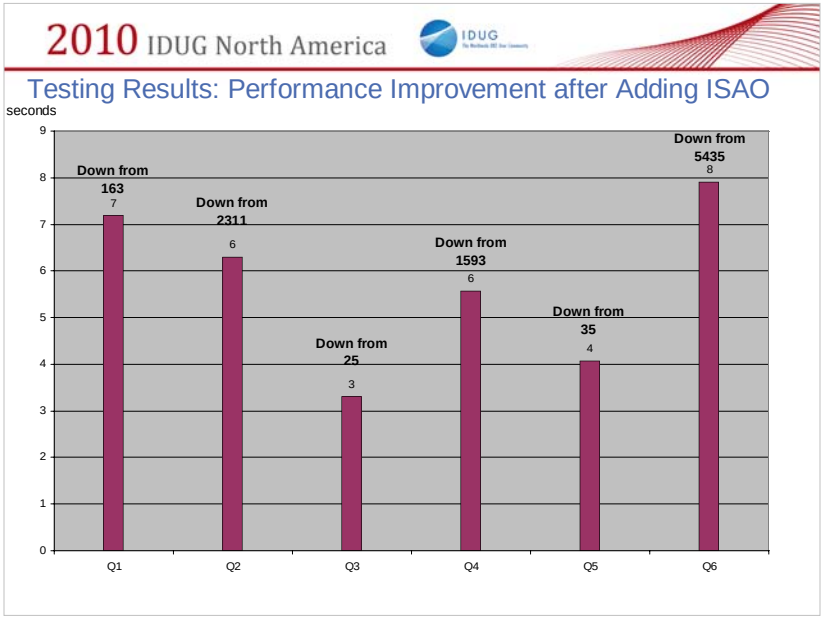

Apart from very significant elapsed time reduction, the kew value of ISAO is relatively small variation among various queries.

Agenda

- Business and Technology Drivers
- Key Design and Operational Features
  - IBM Smart Analytics Optimizer as a virtual DB2 component
  - ISAO engine
  - Marts
  - Query execution
- Supported workloads

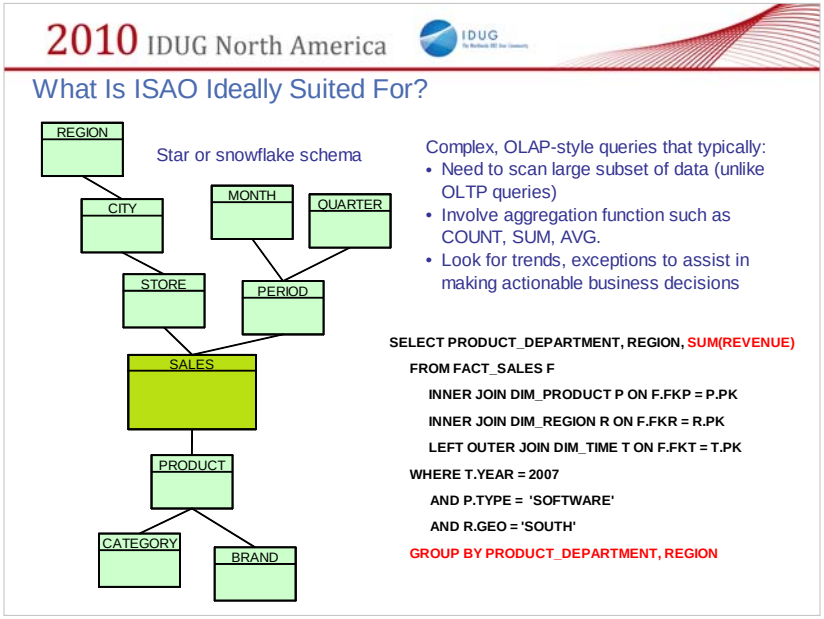

ISAO is not suited for all possible workloads and queries. It's sweet spot are OLAP style queries that access lots of data and include aggregate functions

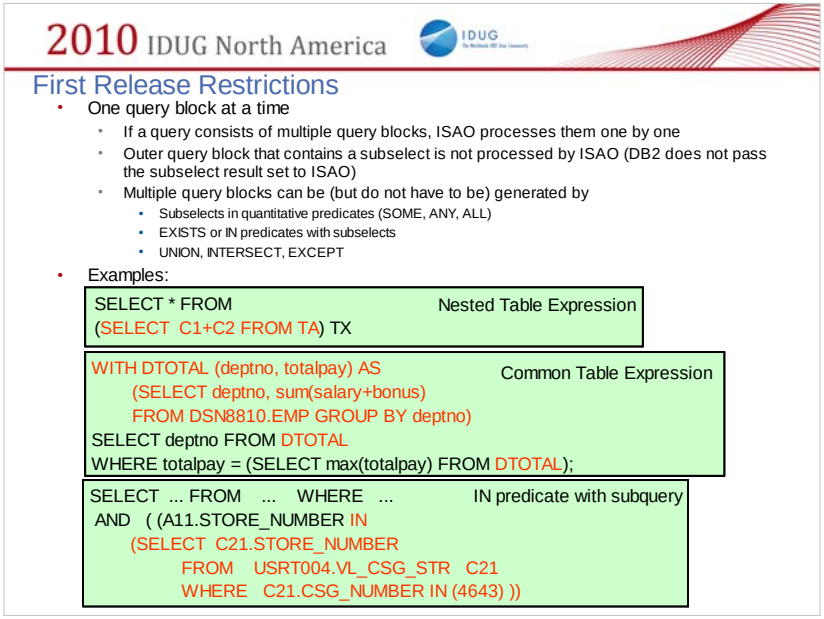

The first release will include some restrictions in terms of which queries and data types are supported and re-routable to ISAO. The following charts list these restrictions.

It's our intention to lift these restrictions in the order of their priority and relevance to our customers.

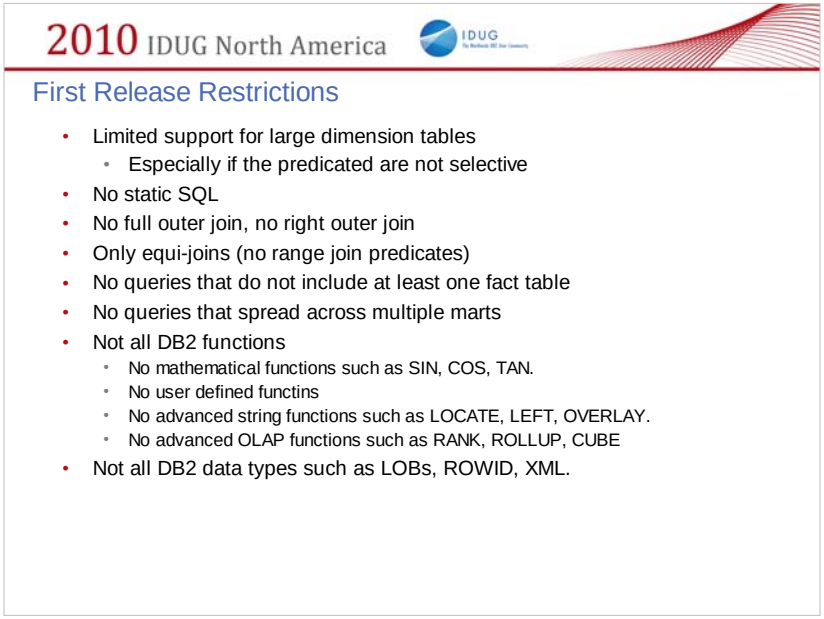

We intend to lift these restrictions in future releases subject to customers needs.

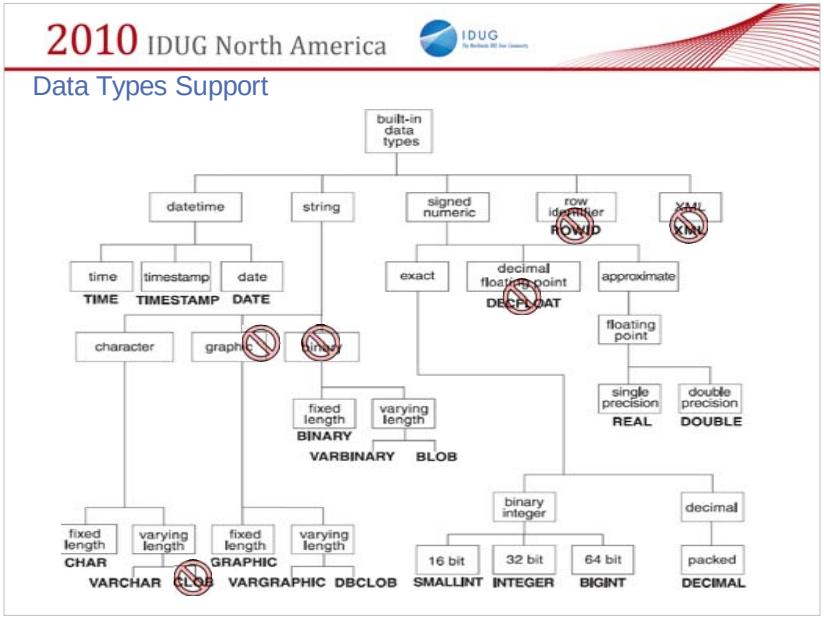

Tables including unsuported data types can still be accelerated through projection.

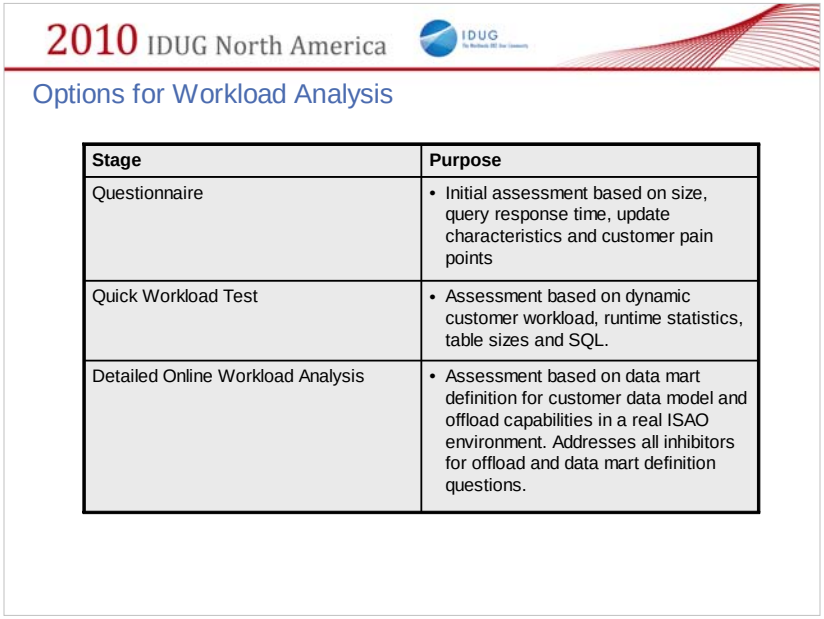

In order to avoid wrong expectations IBM has develop a process for evaluating workload applicability to ISAO. This and the following chart summarize the process.

Quick Workload Test

- Customer
  - Collecting information from dynamic statement cache, supported by step-by-step instruction and REXX script (small effort for customer)
  - Uploading compressed file (up to some MB) to IBM FTP server
- IBM / Center of Excellence
  - Importing data into local database
  - Quick analysis based on known Smart Analytics Optimizer capabilities

Report for a first assessment:
- Query offload potential
- Estimated performance gain
- Estimated CP cost saving

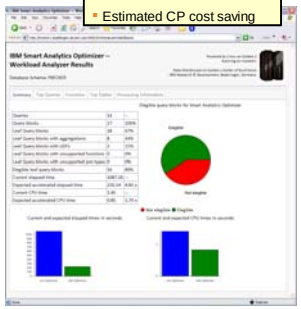

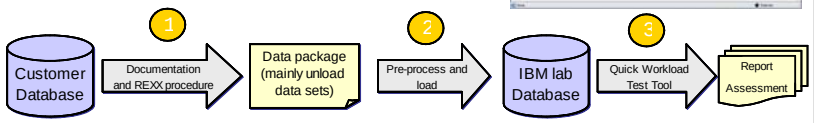

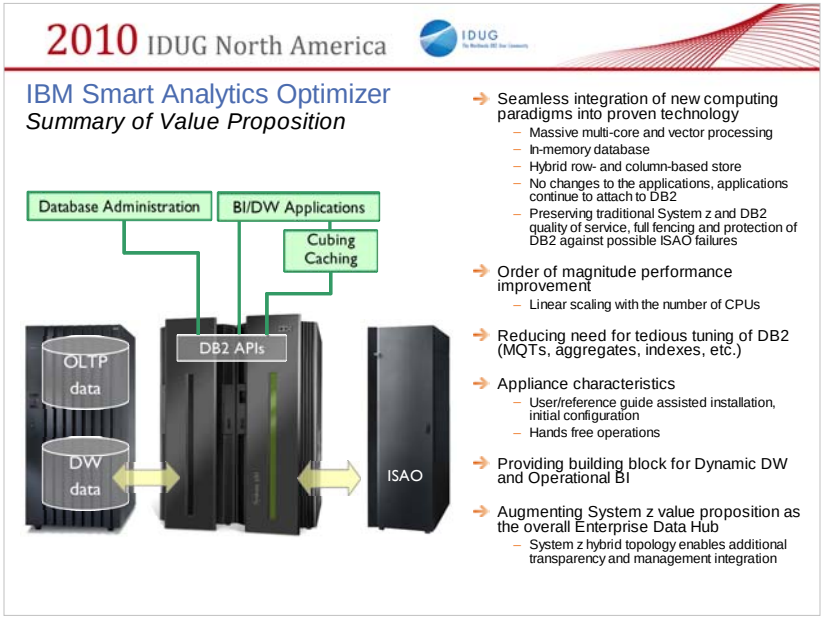

This chart is a one-page summary of the ISAO place in the System z landscape and its main value propositions.

Session Code: A14

Namik Hrle
hrle@de.ibm.com