



IBM Smart Analytics Optimizer for DB2 for z/OS , V1.1 Beta Experiences

William Schray
IBM ATS

08/05/2010 15:00



SHARE in Boston

Disclaimer

© Copyright IBM Corporation [current year]. All rights reserved.

U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM’S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.

IBM, the IBM logo, ibm.com, DB2, and z/OS are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. 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 “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml

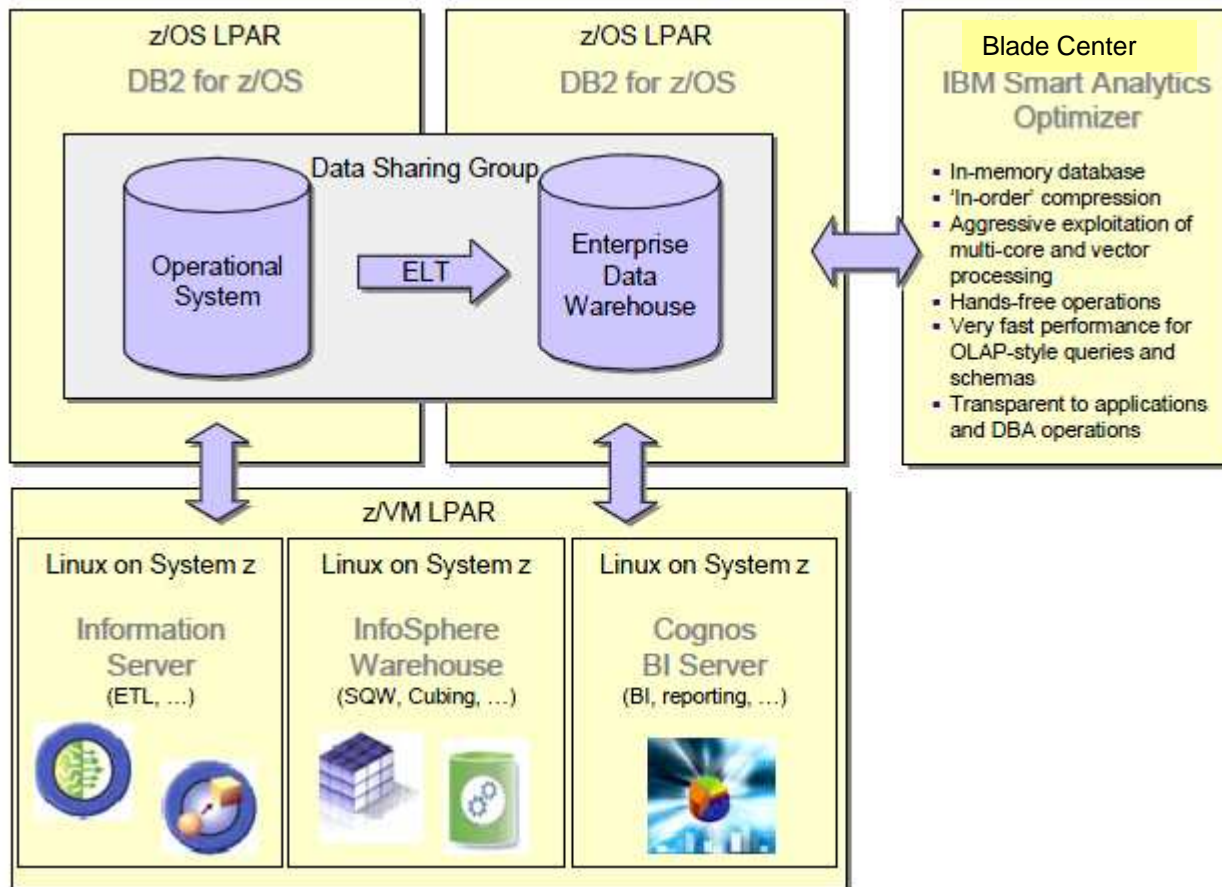
Agenda

- Smart Analytics Optimizer Overview (quickly)
 - Product Positioning
 - Components
 - Type of workloads that qualify
 - Offerings
- Beta Experiences
 - Test databases
 - Determining if queries ran on the accelerator
 - Query performance
 - DB2 data sharing
 - An alternative Data Mart Definition for Workload A and performance measurements

Smart Analytics Optimizer for System z

- An appliance for accelerating certain Data Warehouse and Business Intelligence queries dramatically
- Based on the *Blink* project in IBM's Research Division which 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 Smart Analytics Optimizer design point is to execute queries that are typically found in business intelligence (BI) and data warehousing (DW) applications with fast and predictable response time

Adding the Smart Analytics Optimizer to System z



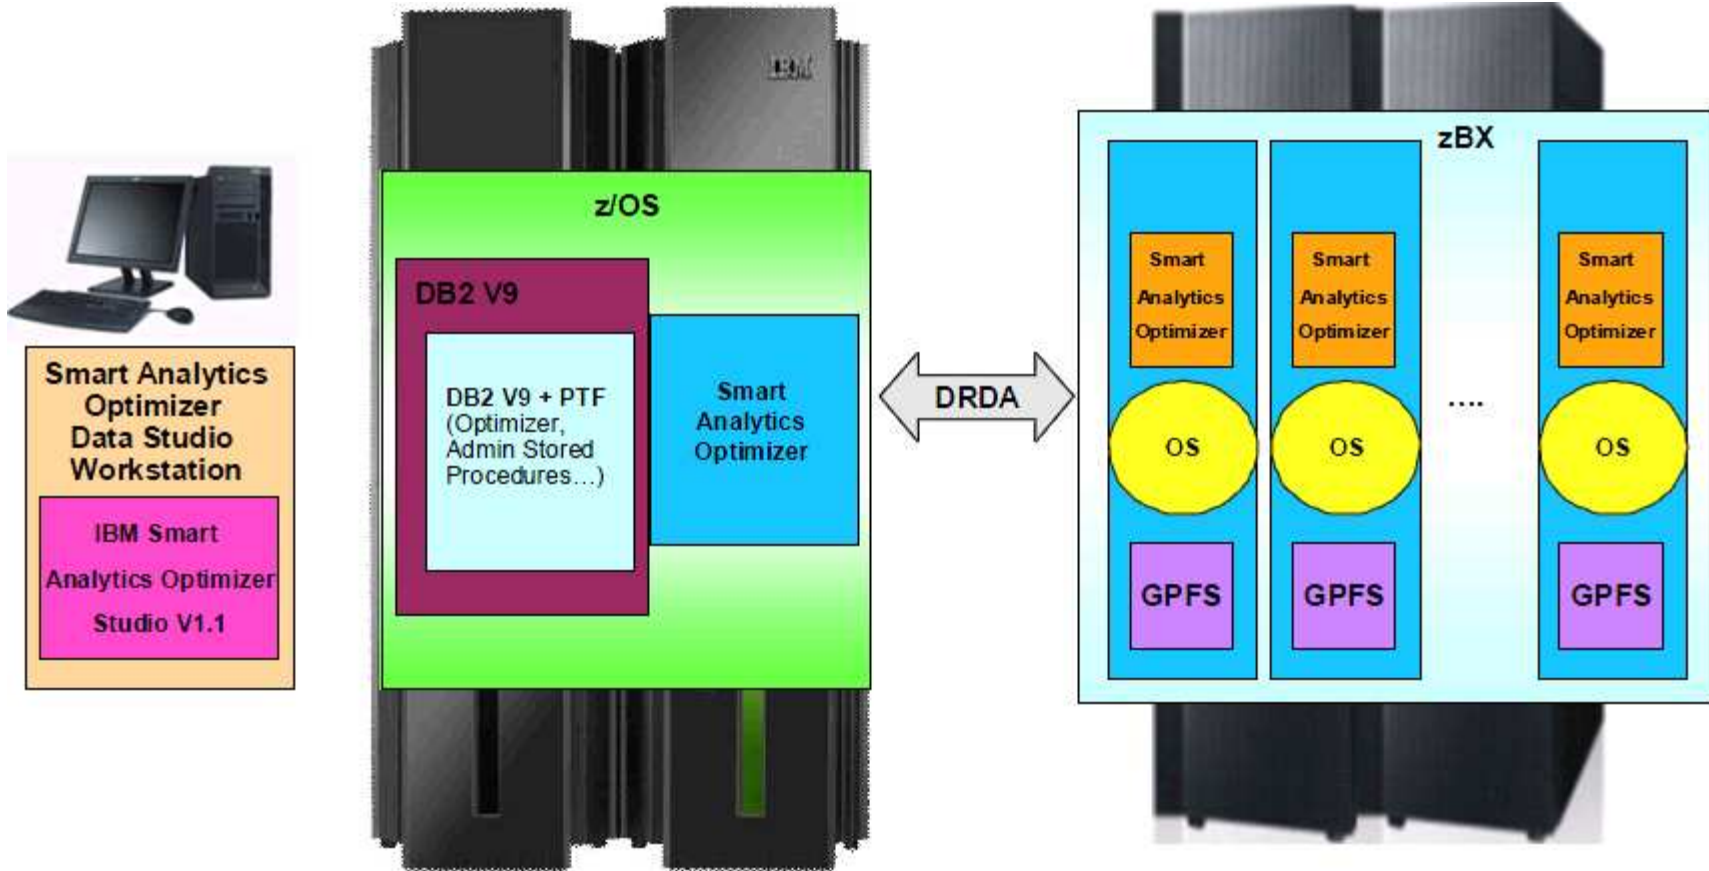
Components of Smart Analytics Optimizer Solution

- IBM Blade center(s) housed in System z BladeCenter Extension (zBX) frame(s) attached to an IBM System z server
 - A closed platform that runs on the blades and includes an operating system and the IBM General Parallel File System V3.x (GPFS™).
- IBM software with installation features for simple initial deployment assisted by graphical tooling. This software is composed of three parts:
 - The unique code developed by IBM and named *Smart Analytics Optimizer Application*.
 - DB2 for z/OS Stored Procedures running on DB2 for z/OS V9 or above with PTF maintenance applied
 - *IBM Smart Analytics Optimizer Studio*. This Eclipse-based GUI runs on a workstation attached to System z and connected to DB2 for z/OS.

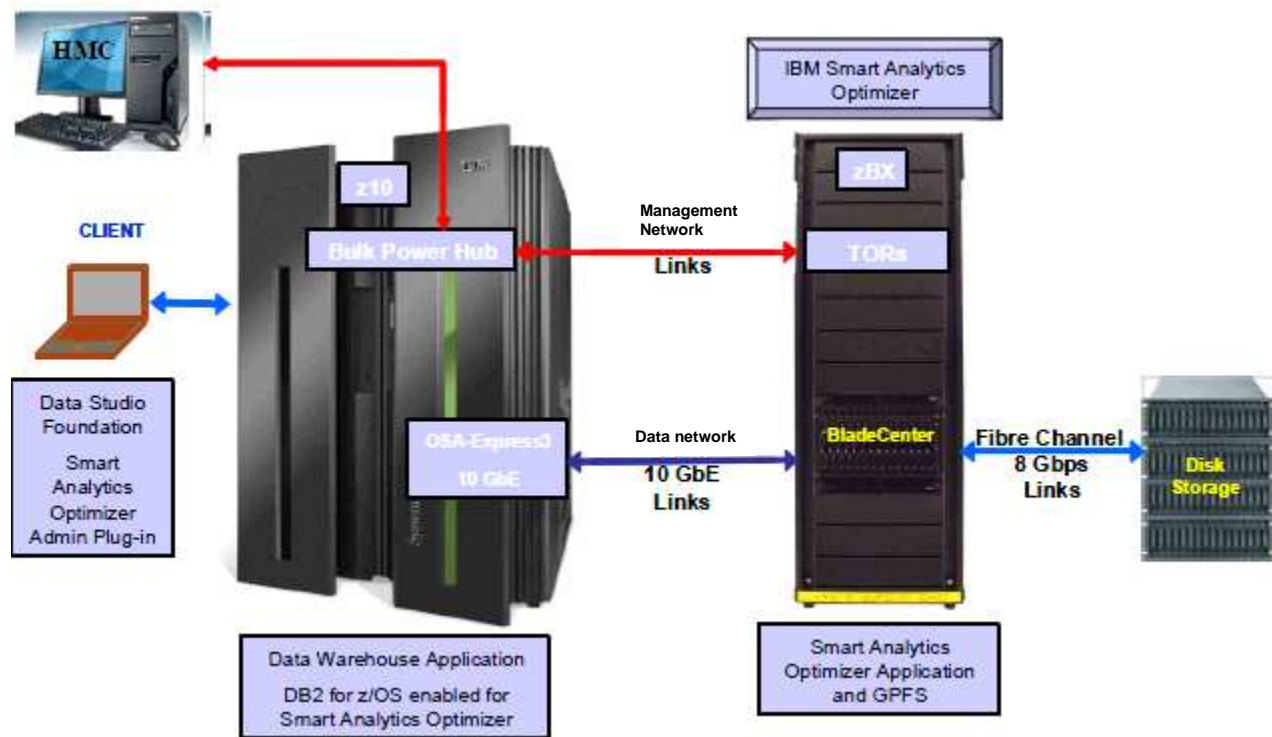
Smart Analytics Optimizer Blade Roles

- The Smart Analytics Optimizer automatically assigns roles to the zBX blades.
 - **Coordinator node** - Accepts and queues work. Distributes and coordinates work for the worker nodes and forwards results to requester (DB2).
 - **Worker node** - Processes the work requests from the coordinator node.

Smart Analytics Optimizer Software Layers



zBX Connectivity



Which types of workload qualify?

- Smart Analytics Optimizer first release requirements for query block offload
 - Dynamic SQL
 - All tables referenced must be defined in one Data Mart
 - **The query must reference at least one fact table in the Data Mart**
 - All columns referenced in the Select list, Where clause, On Clause, Group by clause, Order by Clause, and Having clause must be defined in the Data Mart
 - Join predicates must use =

Which types of workload qualify...?

- Join type must be Inner or Left Outer
- No local filtering predicates in the On clause of Left Outer joins allowed that cannot be pushed down to the Where clause
- Must not contain:
 - User defined functions
 - Mathematical functions (SIN, COS, TAN, EXP, etc.)
 - Character-wise string functions
 - Advanced string functions (LOCATE, LEFT, OVERLAY, POSITION, etc.)
 - Advanced OLAP functions (RANK, DENSE RANK, ROW NUMBER, ROLLUP, CUBE)
 - Certain Data Types (LOB, XML, binary types)

Common Reasons Query Blocks not Offloaded

- Query is statically bound
- CURRENT REFRESH AGE special register set to 0
- Query contains SQL syntax not supported by the accelerator
- Query references a table or column not loaded in a defined accelerator Data Mart
- Query contains Join predicates that do not match defined Data Mart relationships
- Query does not reference at least one fact table
- DB2 optimizer estimates that the query can be executed natively with equal or better performance

Query access path selection and processing

- Accelerator Data Marts are defined in the DB2 on z catalog as one or more Accelerator Query Tables (AQTs)
- The DB2 optimizer's consideration of the AQTs is very similar to the way DB2 considers MQTs, however the internal structure of the AQTs is quite different from the MQTs and is based on BLINK technology
- When AQTs are loaded on the accelerator
 - The fact table(s) are partitioned across the Worker Blades
 - **The dimension tables are replicated in their entirety on each Worker Blade**

Query access path selection and processing...

- The Accelerator processes a single query block in parallel across all of the Worker Blades accessing compressed data in memory on each blade.
 - No access to the attached disk storage occurs during query processing, it is used for start-up and recovery only
 - Certain predicates are processed directly on compressed data in memory, others must de-compress the data in memory before processing the predicate
 - One query block is processed at a time on the accelerator

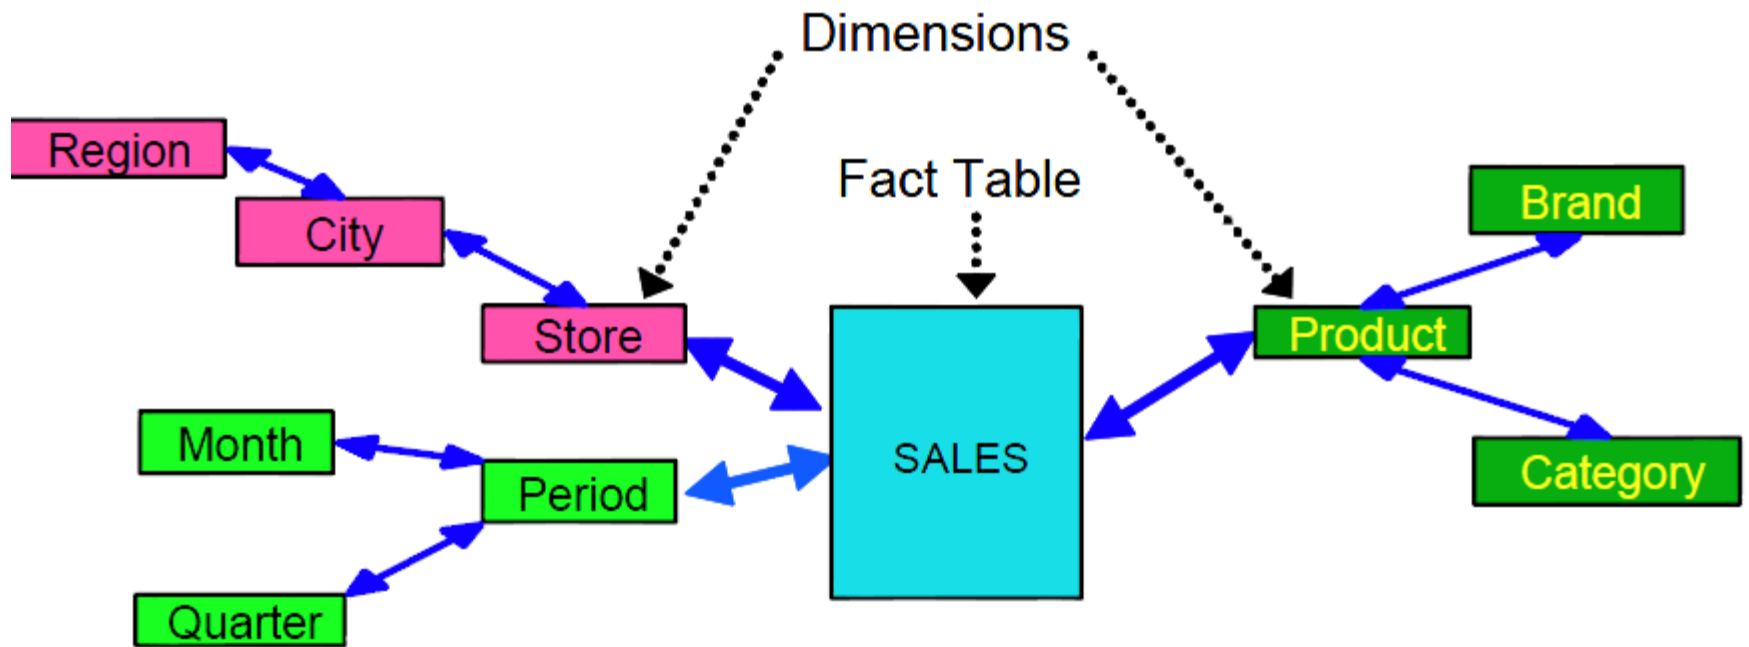
Smart Analytics Optimizer Offerings

- For planning purposes, a 3X compression ratio is a reasonable assumption. In a few cases, 5X was reported but the ratio is highly data dependant.
- Factors to keep in mind when considering required capacity
 - Only the tables and columns used by your longest running queries are required in the accelerator for best performance improvement. However, future flexibility will be enhanced by including more data.
 - Dimension tables are stored in their entirety in each worker blade

Smart Analytics Optimizer Offerings

Solution Offering	Number of Blade Centers	Number of zBX racks	Total Number of blades	Number of Coordinator blade	Number of Worker blades	Blade Fail-over Capability	Total Memory GB (48 GB per blade)	Memory available for Data Mart data GB(32 GB per worker blade)	Memory available for Data Mart data GB with 3X compression
A1-7	1	1	7	2	5	1	336	160	480
A1-14	1	1	14	3	11	2	672	352	1056
A1-28	2	1	28	6	22	5	1344	704	2112
A1-42	3	2	42	9	33	8	2016	1056	3168
A1-56	4	2	56	12	44	11	2688	1408	4224

Star Schema Fact and Dimension Tables



Internal Early Support Program – zBX Model 1

- Objectives
 - Gain experience with Smart Analytics Optimizer
 - Test drive installation process
 - Report problems to ESP sponsors
 - Provide feedback & “customer” perspective to development
- Project team skills
 - z/OS system programmer
 - Hardware and facilities planning
 - DB2 administrator
 - Data warehouse expertise

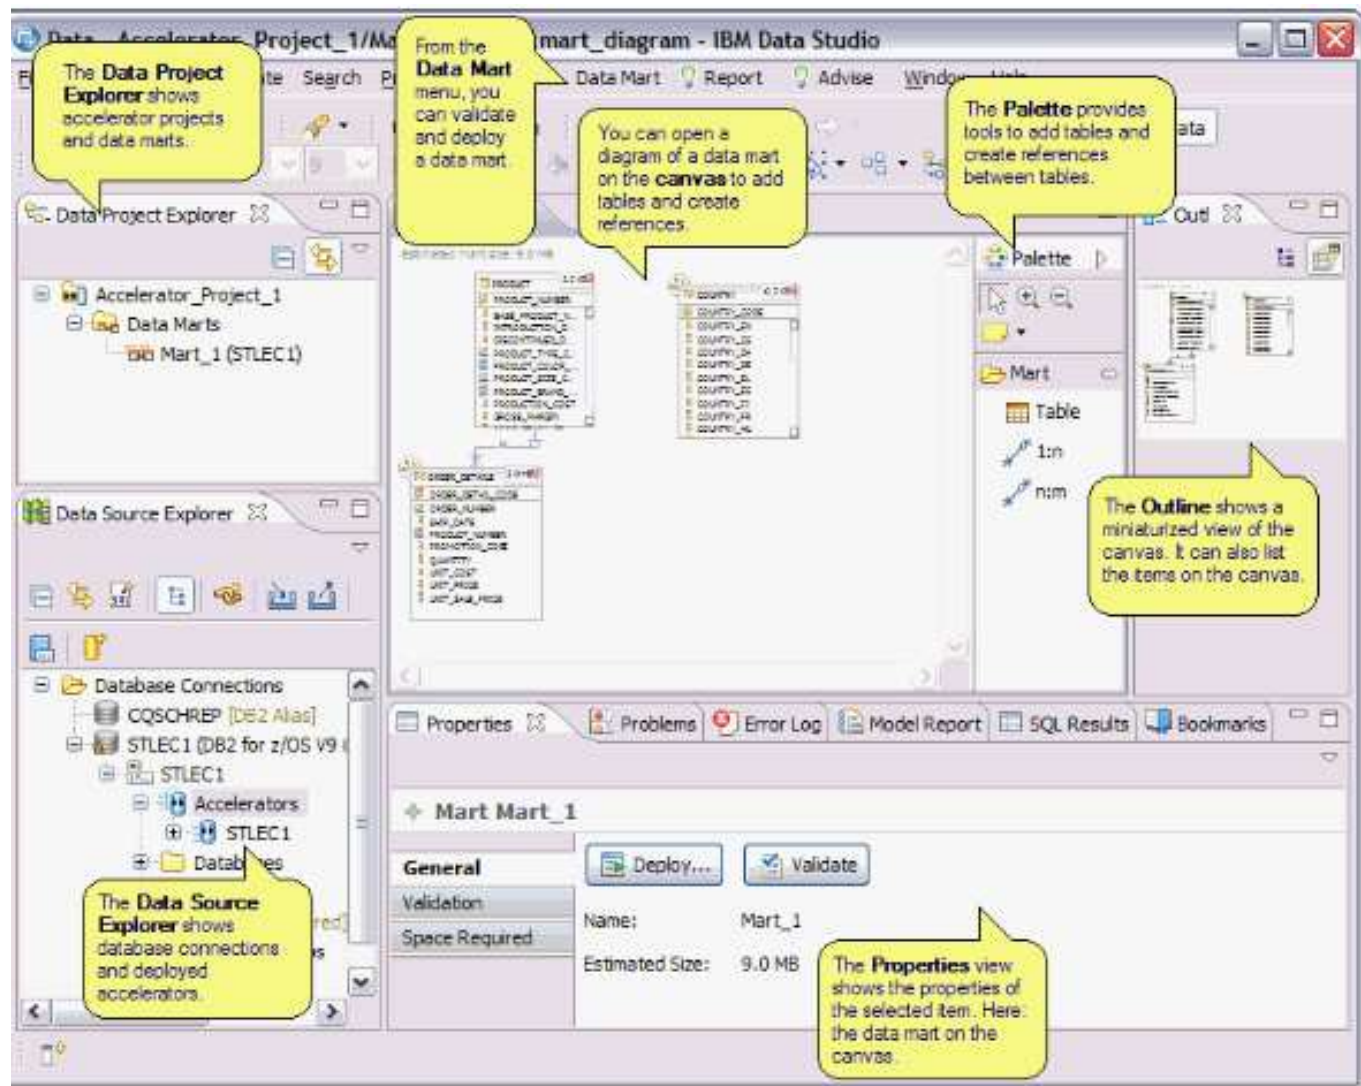
Internal Early Support Program – zBX Model 1

- Status
 - Installed zBX model 1 in February
 - Available to DB2 on two z10 z/OS LPARs
 - Running queries from two workloads on Smart Analytics Optimizer
 - Workload A Lineltem data model
 - Workload B Store_Sales data model

Data Marts

- Data marts created and deployed using Smart Analytics Optimizer Studio
 - Select tables and columns to be in the data mart
- Fact table(s)
 - Data mart must have fact table, usually a large table
 - Partitioned across worker blades
 - Query must reference fact table to be offloaded to optimizer
- Dimension table(s)
 - Each blade contains a copy of the dimension tables
- Table relationships
 - Defined relationships constrains the query joins that can be offloaded to the accelerator

Smart Analytics Optimizer Studio



The screenshot displays the IBM Data Studio Smart Analytics Optimizer Studio interface. The main window shows a data mart diagram with tables and relationships. The interface includes several panes and toolbars:

- Data Project Explorer:** Shows accelerator projects and data marts. Callout: "The Data Project Explorer shows accelerator projects and data marts."
- Data Source Explorer:** Shows database connections and deployed accelerators. Callout: "The Data Source Explorer shows database connections and deployed accelerators."
- Diagram Canvas:** The central area for creating and editing data mart diagrams. Callout: "From the Data Mart menu, you can validate and deploy a data mart." and "You can open a diagram of a data mart on the canvas to add tables and create references."
- Palette:** Provides tools to add tables and create references between tables. Callout: "The Palette provides tools to add tables and create references between tables."
- Outline:** Shows a miniaturized view of the canvas and lists the items on it. Callout: "The Outline shows a miniaturized view of the canvas. It can also list the items on the canvas."
- Properties View:** Shows the properties of the selected item (the data mart). Callout: "The Properties view shows the properties of the selected item. Here, the data mart on the canvas." The Properties view for "Mart Mart_1" includes:

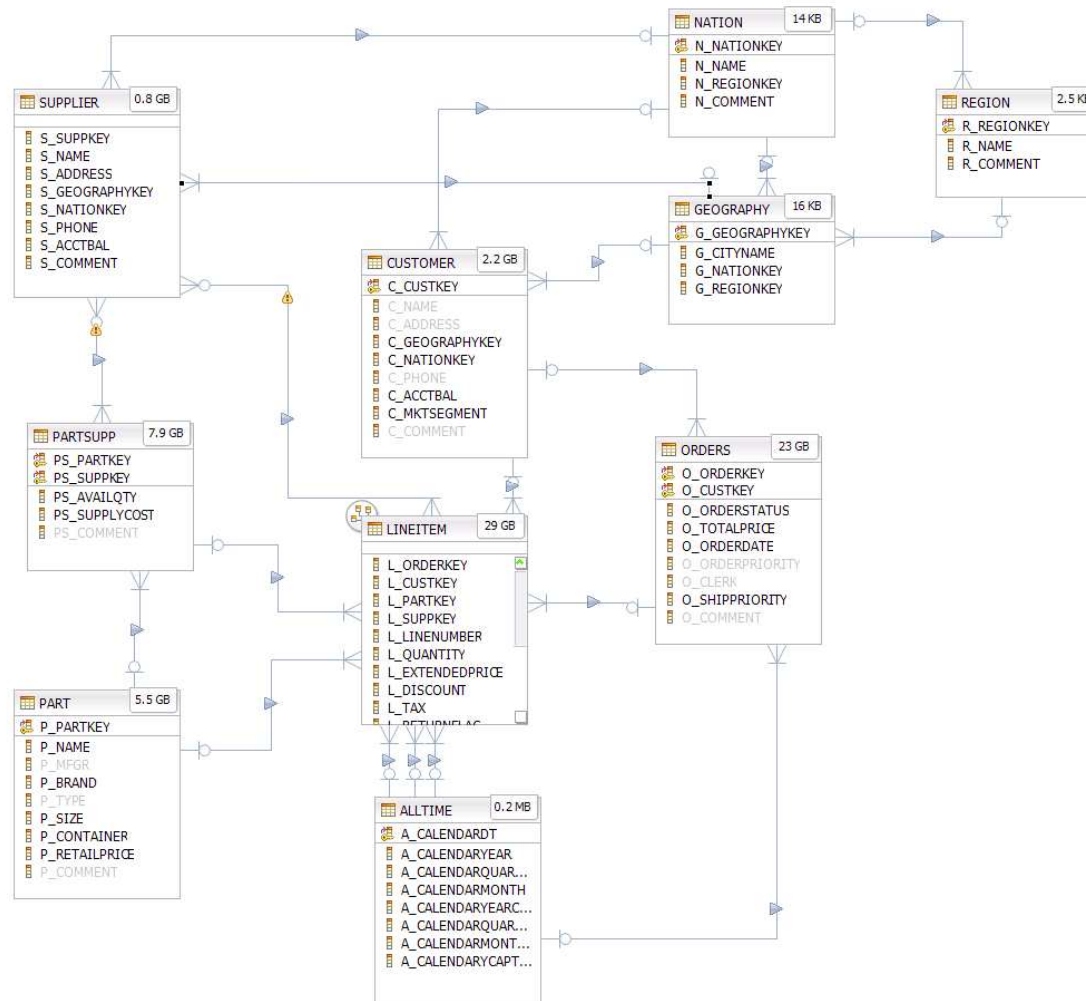
Mart Mart_1	
General	
[Deploy...] [Validate]	
Validation	
Space Required	Name: Mart_1
	Estimated Size: 9.0 MB

Workloads Used During Beta Testing

- Two different query workloads were created to test IBM Smart Analytics Optimizer acceleration capabilities
 - Workload A based on a 100GB version of the TPC-H data model
 - Workload B based on a 100GB version of the TPC-DS data model
- The query workloads used do not contain any of the queries typically used with these data models.
- These performance tests were not official benchmarks and they did not adhere to the official benchmarks rules, nor have we submitted results to any formal review by the TPC board. We have merely constructed workloads that resemble, but are not the same as, the industry standard TPC-H benchmarks.
- The query workloads will hereafter be referred to as 'workload A' and 'workload B'.

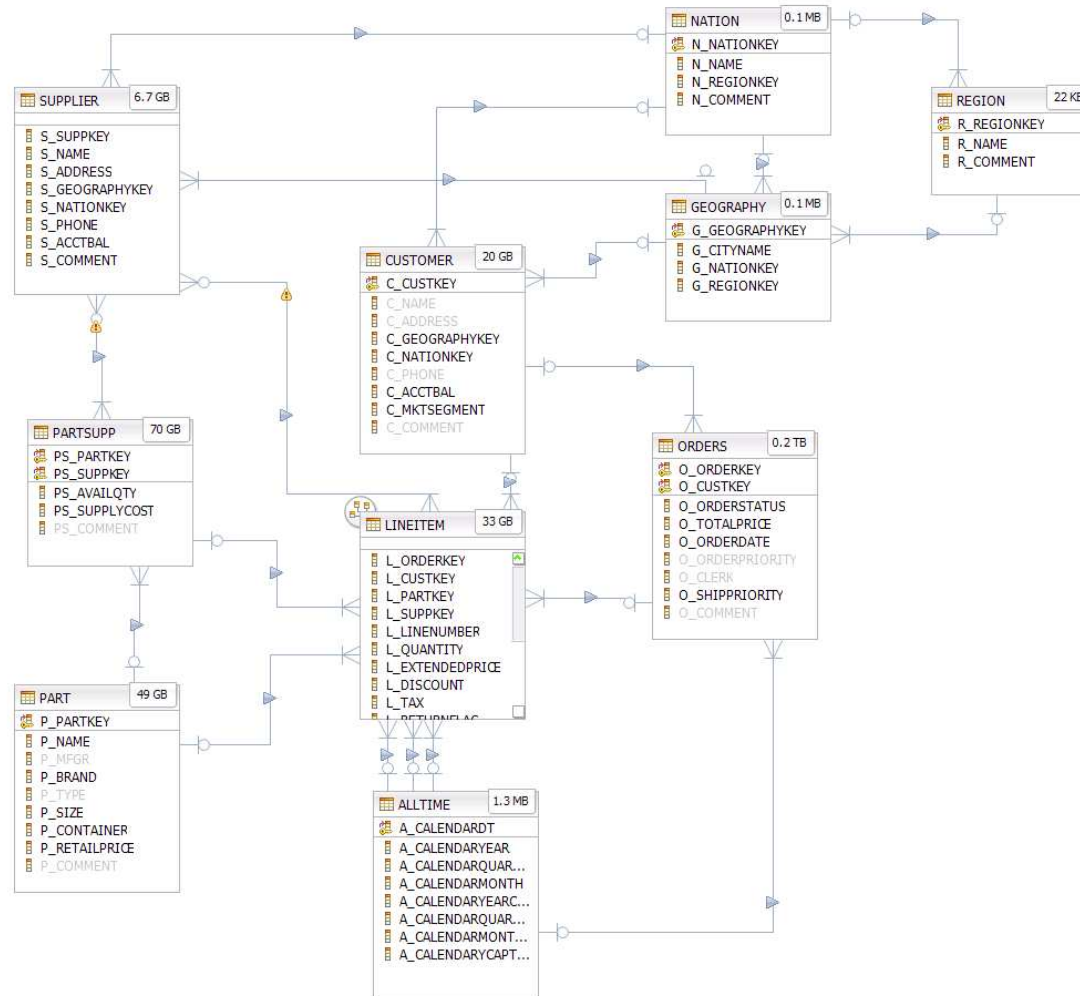
Workload A Lineltem Data Mart – Space estimate for 7-blades

Estimated size: 68 GB



Workload A Lineltem Data Mart – Space estimate for 56-blades

Estimated size: 0.4 TB



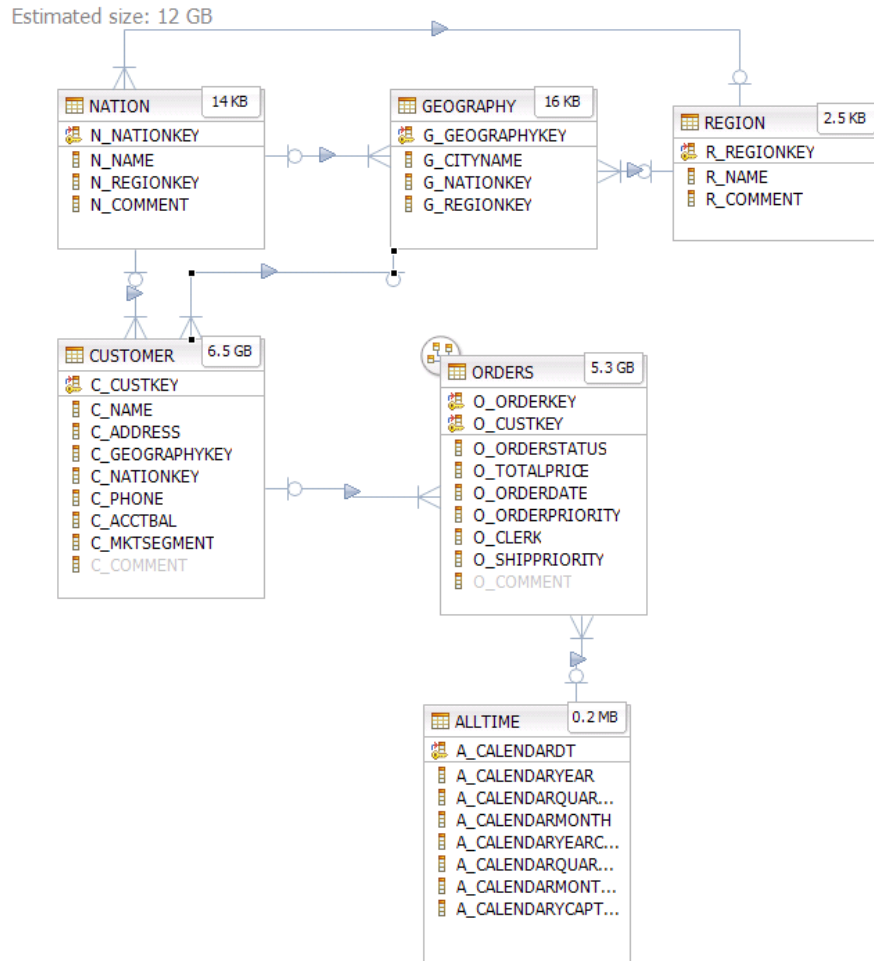
Workload A Data Mart(s)

- The Workload A data model is not a star schema, but rather a relational schema like you find in an ODS
- There are many of the queries in my test workload that would not offload to the accelerator because they did not reference the Lineitem table, the only fact table in this Data Mart definition.
- To allow all of my long running queries to offload to the accelerator, I needed to define the following additional data marts which are shown on subsequent pages:
 - Workload A Order data mart
 - Workload A PartSupp data mart
 - Workload A Supplier data mart
- When this presentation was being reviewed, the lab suggested an alternative data mart definition in which tables Lineitem, Order, PartSupp, Supplier, and Part are defined as fact tables is a single data mart definition.

Workload A query Joins that offloaded to LINEITEM data mart

- A01 – LINEITEM
- A13 – LINEITEM, SUPPLIER, GEOGRAPHY
- A14 – LINEITEM, SUPPLIER, GEOGRAPHY
- A15 – LINEITEM, SUPPLIER, GEOGRAPHY, NATION
- A16 – LINEITEM, PART
- A17 – LINEITEM, PART
- A18 – LINEITEM, ORDERS (Run time problem)
- A19 – LINEITEM, CUSTOMER, GEOGRAPHY
- A20 – LINEITEM, CUSTOMER, NATION
- A21 – LINEITEM, CUSTOMER, NATION
- A24 – LINEITEM, ALLTIME
- A25 – LINEITEM, ALLTIME

Workload A Lineltem Order Data Mart – Space estimate for 7-blades

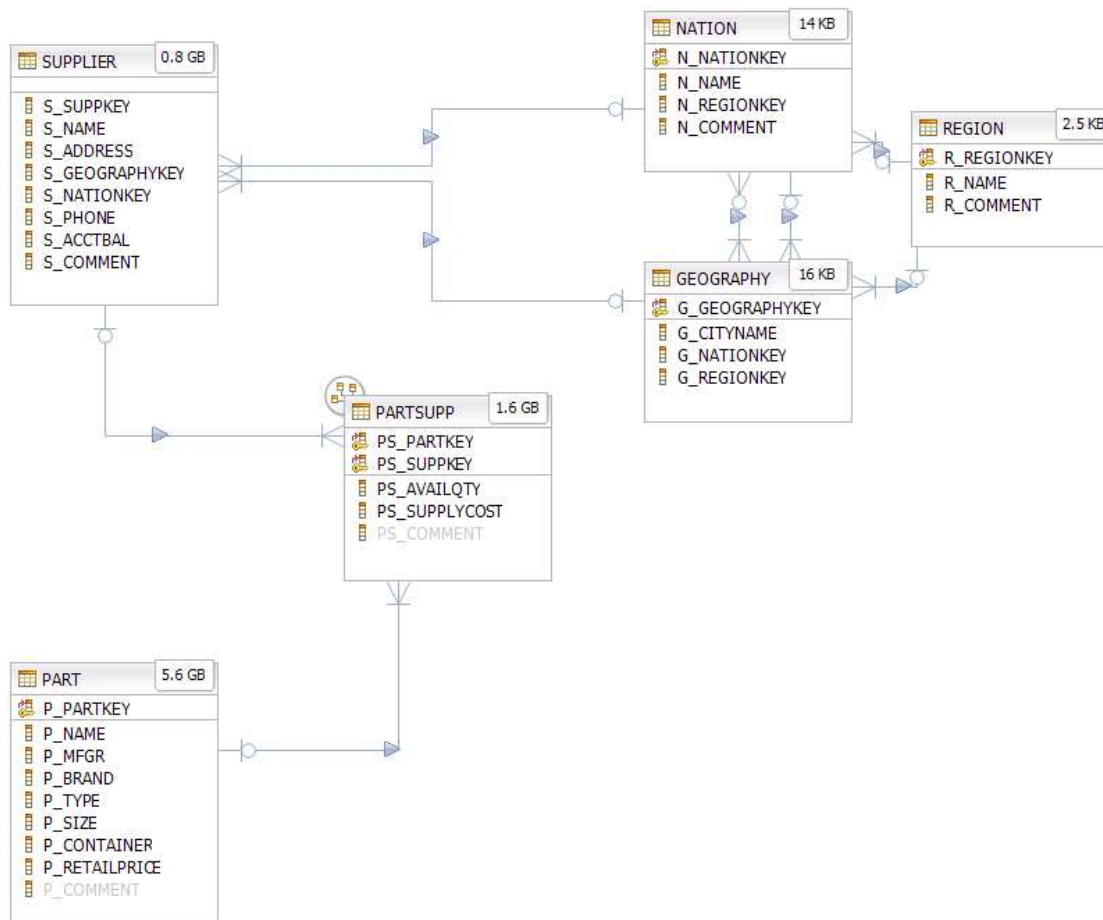


Workload A query Joins that offloaded to ORDER data mart

- A04 – ORDERS
- A05 – ORDERS, CUSTOMER
- A06 – ORDERS, CUSTOMER, NATION
- A07 – ORDERS, CUSTOMER
- A09 – ORDERS, CUSTOMER, GEOGRAPHY
- A10 – ORDERS, CUSTOMER
- A23 – ORDERS, ALLTIME
- A26 – ORDERS, CUSTOMER, NATION, REGION

Workload A Lineltem PartSupp Data Mart – Space estimate for 7-blades

Estimated size: 8.0 GB



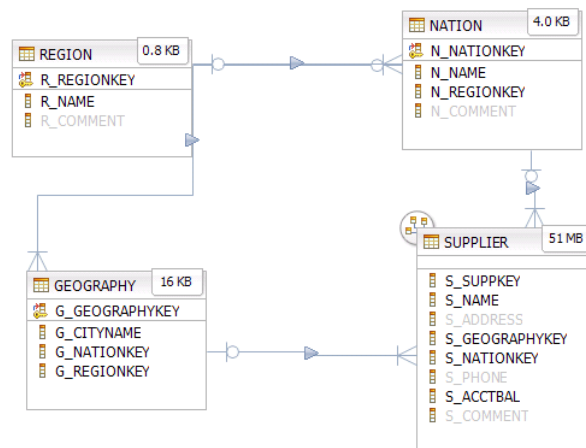
Workload A query Joins that offloaded to PARTSUPP data mart



- A22 – PARTSUPP, SUPPLIER

Workload A Lineltem Supplier Data Mart – Space estimate for 7-blades

Estimated size: 51 MB



Workload A query Joins that offloaded to SUPPLIER data mart



- A11 – SUPPLIER, GEOGRAPHY
- A12 – SUPPLIER, GEOGRAPHY

Workload A Remaining query Joins that did not offload

- A02 – CUSTOMER
- A03 – PART
- These queries probably could have been made to offload with additional data mart definitions but these queries ran so fast on z that there seemed to be no point in adding additional data marts.

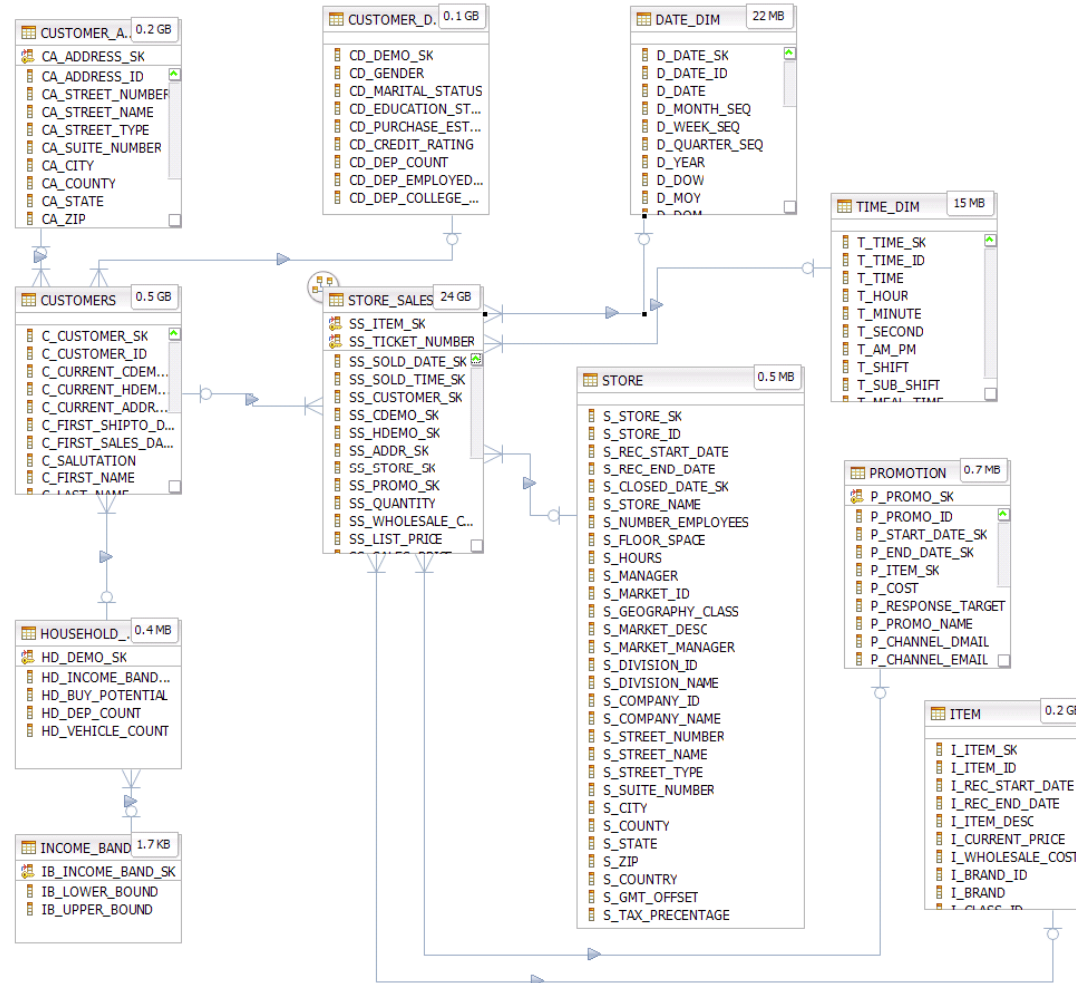
Workload B Store_Sales Data Mart

- The Workload B Store_Sales data model is a classic star schema which allows all of the 13 queries in my test workload to offload using a single Data Mart definition.
- Since the dimension tables only represent 1.3 % of the data in the data model, this data mart scales fairly linearly as blade centers were added to the accelerator

Workload B Store_Sales Data Mart – Space estimate for 7-blades

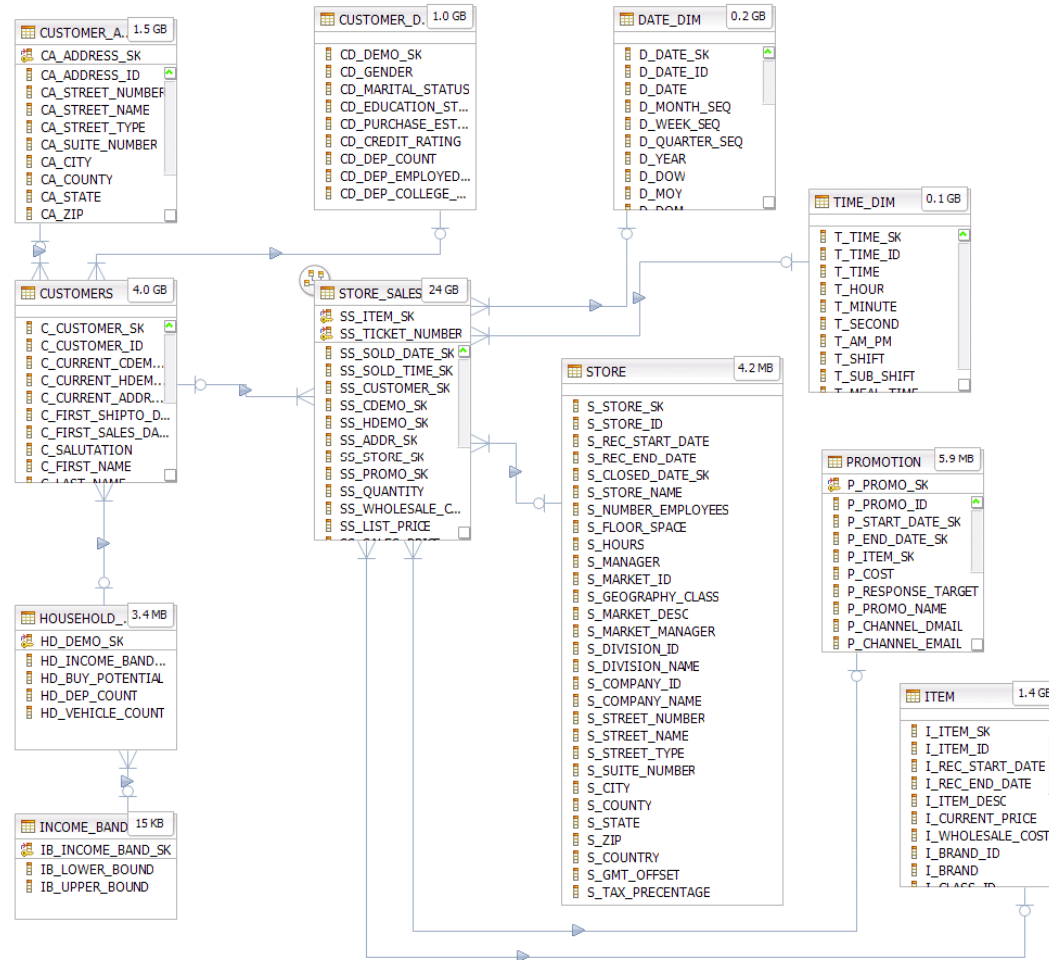


Estimated size: 25 GB



Workload B Store_Sales Data Mart – Space estimate for 56-blades

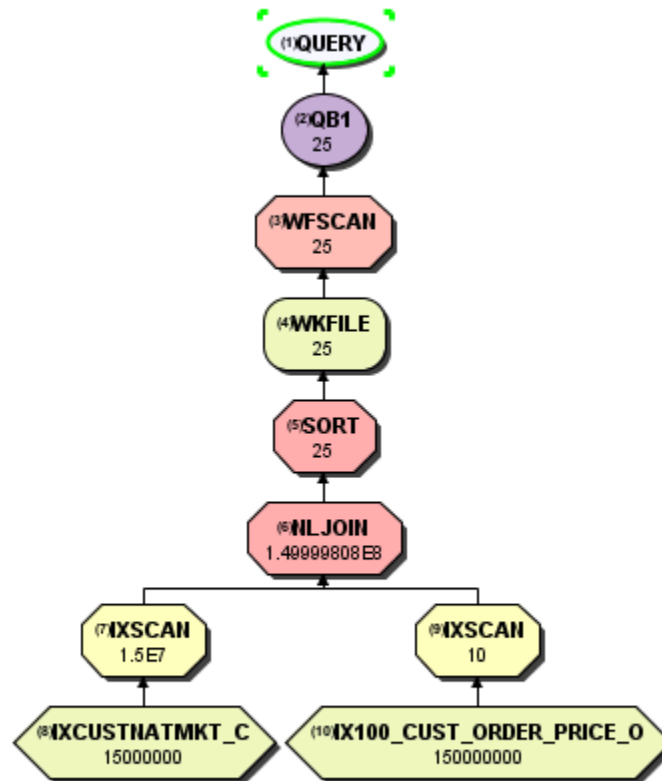
Estimated size: 32 GB




Determining if Queries will be offloaded

- Run DB2 EXPLAIN for the query
 - PLAN_TABLE refers to AQT table(s)
 - DSN_QUERYINFO_TABLE lists reason not eligible for offload (new EXPLAIN table)
- Accelerator internal errors reported in DB2 MSTR message log and the system log
 - Detailed error message is written to IFCID 191

Access Plan Graph with 'SET CURRENT REFRESH AGE 0:'



Access Plan Graph with 'SET CURRENT REFRESH AGE ANY;'





OMPE Accounting Report with Accelerator Offload

ELAPSED TIME DISTRIBUTION

```

-----
APPL |
DB2  |
SUSP |===== > 100%

```

CLASS 2 TIME DISTRIBUTION

```

-----
CPU  |
NOTACC |
SUSP |===== > 100%

```

TIMES/EVENTS	APPL(CL.1)	DB2 (CL.2)	IFI (CL.5)	CLASS 3 SUSPENSIONS	ELAPSED TIME	EVENTS	HIGHLIGHTS
ELAPSED TIME	31.086663	31.082156	N/P	LOCK/LATCH(DB2+IRLM)	0.000000	0	THREAD TYPE : ALLIED
NONNESTED	31.086663	31.082156	N/A	SYNCHRON. I/O	0.000000	0	TERM.CONDITION: NORMAL
STORED PROC	0.000000	0.000000	N/A	DATABASE I/O	0.000000	0	INVOKE REASON : DEALLOC
UDF	0.000000	0.000000	N/A	LOG WRITE I/O	0.000000	0	COMMITTS : 1
TRIGGER	0.000000	0.000000	N/A	OTHER READ I/O	0.000000	0	ROLLBACK : 0
				OTHER WRTE I/O	0.000000	0	SVPT REQUESTS : 0
CP CPU TIME	0.003298	0.001531	N/P	SER.TASK SWTCH	31.079452	14	SVPT RELEASE : 0
AGENT	0.003298	0.001531	N/A	UPDATE COMMIT	0.000000	0	SVPT ROLLBACK : 0
NONNESTED	0.003298	0.001531	N/P	OPEN/CLOSE	0.000000	0	INCREM.BINDS : 0
STORED PROC	0.000000	0.000000	N/A	SYSLGRNG REC	0.000000	0	UPDATE/COMMIT : 0.00
UDF	0.000000	0.000000	N/A	EXT/DEL/DEF	0.000000	0	SYNCH I/O AVG. : N/C
TRIGGER	0.000000	0.000000	N/A	OTHER SERVICE	31.079452	14	PROGRAMS : 1
PAR.TASKS	0.000000	0.000000	N/A	ARC.LOG(QUIES)	0.000000	0	MAX CASCADE : 0
				LOG READ	0.000000	0	PARALLELISM : NO
IIPCP CPU	0.000000	N/A	N/A	DRAIN LOCK	0.000000	0	
				CLAIM RELEASE	0.000000	0	
IIP CPU TIME	0.000000	0.000000	N/A	PAGE LATCH	0.000000	0	
STORED PROC	0.000000	0.000000	N/A	NOTIFY MSGS	0.000000	0	
				GLOBAL CONTENTION	0.000000	0	
SUSPEND TIME	0.000000	31.079452	N/A	COMMIT PH1 WRITE I/O	0.000000	0	
AGENT	N/A	31.079452	N/A	ASYNCH CF REQUESTS	0.000000	0	
PAR.TASKS	N/A	0.000000	N/A	TCP/IP LOB	0.000000	0	
STORED PROC	0.000000	N/A	N/A	TOTAL CLASS 3	31.079452	14	
UDF	0.000000	N/A	N/A				
NOT ACCOUNT.	N/A	0.001173	N/A				
DB2 ENT/EXIT	N/A	52	N/A				
EN/EX-STPROC	N/A	0	N/A				
EN/EX-UDF	N/A	0	N/A				
DCAPT.DESCR.	N/A	N/A	N/P				
LOG EXTRACT.	N/A	N/A	N/P				

Smart Analytics Optimizer Instrumentation

- DB2 commands to start, stop, and show status
 - -DSNA Start Accel
 - -DSNA Stop Accel
 - -DSNA Display Accel(*) [Detail]
- Instrumentation provided by Accelerator
 - DB2 SMF 100 and 101 records
 - Accelerator traces started in Smart Analytics Optimizer Studio

Output from “/-DSNA Display Accel(*) Detail”

```

NC0000000 SYSA      10213 07:30:37.76 SCHRAY   00000210  -DSNA DIS ACCEL(*) DETAIL
NR0000000 SYSA      10213 07:30:37.76 STC16212 00000010  DSNX810I  -DSNA DSNX8CMD DISPLAY ACCEL FOLLOWS -
MR0000000 SYSA      10213 07:30:37.76 STC16212 00000010  DSNX830I  -DSNA DSNX8CDA 689
DR                  689 00000010  ACCELERATOR                                MEMB  STATUS  REQUESTS  QUED  MAXQ  FAIL
DR                  689 00000010  -----
DR                  689 00000010  ZBX1                                         DSNX  STARTED  65      0    1    0
DR                  689 00000010  PORT=1400 IPNAME=07C2C17D
DR                  689 00000010  IPADDR=10.9.130.1 HEALTHY
DR                  689 00000010  IPADDR=10.9.130.2 HEALTHY
DR                  689 00000010  DETAIL STATISTICS
DR                  689 00000010    LEVEL = AQT01010
DR                  689 00000010    STATUS = FULLY OPERATIONAL
DR                  689 00000010    AVERAGE QUEUE WAIT                                =      0
DR                  689 00000010    MAXIMUM QUEUE WAIT                                =      1
DR                  689 00000010    PROCESSING CAPACITY                                = 657139
DR                  689 00000010    TOTAL NUMBER OF PROCESSORS                         =    112
DR                  689 00000010    AVERAGE CPU UTILIZATION ON COORDINATOR NODES     =    .62%
DR                  689 00000010    NUMBER OF ACTIVE COORDINATOR NODES                 =      2
DR                  689 00000010    PHYSICAL MEMORY AVAILABLE ON COORDINATOR NODES    = 47180 MB
DR                  689 00000010    PHYSICAL MEMORY AVERAGE USAGE ON COORDINATOR NODES = 1094 MB
DR                  689 00000010    AVERAGE CPU UTILIZATION ON WORKER NODES          =    .62%
DR                  689 00000010    NUMBER OF ACTIVE WORKER NODES                       =      5
DR                  689 00000010    PHYSICAL MEMORY AVAILABLE ON WORKER NODES          = 39864 MB
DR                  689 00000010    PHYSICAL MEMORY AVERAGE USAGE ON WORKER NODES    = 8410 MB
DR                  689 00000010    SHARED MEMORY DATA AVAILABLE ON WORKER NODES     = 4290 MB
DR                  689 00000010    SHARED MEMORY DATA AVERAGE USAGE ON WORKER NODES = 27861 MB
DR                  689 00000010    MAXIMUM SHARED MEMORY DATA IN USE ON WORKER NODES = 27873 MB
ER                  689 00000010  DISPLAY ACCEL REPORT COMPLETE
NR0000000 SYSA      10213 07:30:37.76 STC16212 00000010  DSN9022I  -DSNA DSNX8CMD '-DISPLAY ACCEL' NORMAL COMPLETION

```

After Workload A execution



```

NR0000000 SYSA      10213 07:43:24.34 STC16212 00000010 DSNX810I  -DSNA DSNX8CMD DISPLAY ACCEL FOLLOWS -
MR0000000 SYSA      10213 07:43:24.34 STC16212 00000010 DSNX830I  -DSNA DSNX8CDA 821
DR                821 00000010 ACCELERATOR                                MEMB  STATUS  REQUESTS  QUED  MAXQ  FAIL
DR                821 00000010 -----
DR                821 00000010 ZBX1                                        DSNX  STARTED                                89    0    1    0
DR                821 00000010 PORT=1400 IPNAME=07C2C17D
DR                821 00000010 IPADDR=10.9.130.1 HEALTHY
DR                821 00000010 IPADDR=10.9.130.2 HEALTHY
DR                821 00000010 DETAIL STATISTICS
DR                821 00000010   LEVEL   = AQT01010
DR                821 00000010   STATUS   = FULLY OPERATIONAL
DR                821 00000010   AVERAGE QUEUE WAIT                               =          0
DR                821 00000010   MAXIMUM QUEUE WAIT                               =          1
DR                821 00000010   PROCESSING CAPACITY                               =      657139
DR                821 00000010   TOTAL NUMBER OF PROCESSORS                       =         112
DR                821 00000010   AVERAGE CPU UTILIZATION ON COORDINATOR NODES    =          .00%
DR                821 00000010   NUMBER OF ACTIVE COORDINATOR NODES               =          2
DR                821 00000010   PHYSICAL MEMORY AVAILABLE ON COORDINATOR NODES  =      47181 MB
DR                821 00000010   PHYSICAL MEMORY AVERAGE USAGE ON COORDINATOR NODES =      1093 MB
DR                821 00000010   AVERAGE CPU UTILIZATION ON WORKER NODES         =          .00%
DR                821 00000010   NUMBER OF ACTIVE WORKER NODES                    =          5
DR                821 00000010   PHYSICAL MEMORY AVAILABLE ON WORKER NODES        =      39864 MB
DR                821 00000010   PHYSICAL MEMORY AVERAGE USAGE ON WORKER NODES  =       8410 MB
DR                821 00000010   SHARED MEMORY DATA AVAILABLE ON WORKER NODES   =       4290 MB
DR                821 00000010   SHARED MEMORY DATA AVERAGE USAGE ON WORKER NODES =      27861 MB
DR                821 00000010   MAXIMUM SHARED MEMORY DATA IN USE ON WORKER NODES =      27873 MB
ER                821 00000010 DISPLAY ACCEL REPORT COMPLETE
NR0000000 SYSA      10213 07:43:24.34 STC16212 00000010 DSN9022I  -DSNA DSNX8CMD '-DISPLAY ACCEL' NORMAL COMPLETION

```

How Can We Determine How Many Queries Executed Successfully on the Accelerator?

- From the previous two accelerator displays, we can tell that 24 of the 26 (89 – 65) of the Workload A queries completed successfully on the accelerator
- For each individual job step, we could determine from an OMPE trace which of the queries ran successfully on the accelerator
- When queries are expected to run on the accelerator from EXPLAIN results but do not successfully run on the accelerator, how do we tell what internal errors occurred?

Smart Analytics Optimizer Internal Error Debugging

System LOG error message (also found in the DB2 MSTR started task)

```
DSNX871I -DSNA DSNX8DJN DRDA EXCEPTION CONDITION IN 489
RESPONSE FROM ACCELERATOR=ZBX1 FOR THREAD WITH
LUWID=USIBMWZ.DSNAAPPL.C6364562E46A=0
REASON=00E7000A
ERROR ID=DSNX8OPQ0005
CORRELATION ID=TPCHQ01
CONNECTION ID=BATCH
IFCID=0191
SEE TRACE RECORD WITH IFCID SEQUENCE NUMBER=00000077
```

OMPE RECTRACE of IFCID 191

```
SCHRAY  BATCH      C63633672040 'BLANK'          'BLANK'          'BLANK'
SCHRAY  TCPDSQ10 TSO          08:51:26.89095159 171408    3 191 DDM LEVEL 6B   NETWORKID:  USIBMWZ    LUNAME:  DSNAAPPL  LUWSEQ:    1
DSNTEP91 'BLANK'          N/P                                OBJECTS
```

```
|0000 00F50000 C4D9C4C1 F0F1F9F1 E9C2E7F1 40404040 40404040 40404040 40400001 |.5..DRDA0191ZBX1 ..
|0020 00000000 00000000 00000000 00E7F803 C4E2D5E7 F8D9D7C1 0071C4E2 D5E7F8D9 |.....X8.DSNX8RPA..DSNX8R
|0040 D7C1F0F0 F7F1200D D7F10000 00000000 00000000 00000000 00000000 00000000 |PA0071..P1.....
|0060 00040202 05060600 02010000 00000000 00000000 00000000 0000C4D9 C4C1D9D7 |.....DRDARP
|0080 D3E8C4D9 C4C1E2E4 C3C31233 00010000 00AC0000 00000000 00000000 00000000 |LYDRDASUCC.....
|00A0 00000000 00000000 00000000 0000C4D9 C4C1D9D7 E840003F D0020005 00391233 |.....DRDARPY ..}.....
|00C0 002F1153 E3888540 81838385 93859981 A3969940 83819540 9596A340 88819584 |...The accelerator can not hand
|00E0 93854094 96998540 98A48599 8985A200 06114900 08                                |le more queries.....
```

Performance Measurements

- The Smart Analytics Optimizer Beta performance measurements that follow should be considered in the following context:
 - The actual performance that any user will experience will vary depending upon considerations such as the amount of data on each blade in the AQT table being referenced. Therefore, no assurance can be given that an individual user will achieve response time improvements equivalent to those stated here.
 - Queries built for the two data models to exercise the Smart Analytics Optimizer during Beta testing were designed to be likely to offload to the accelerator. The queries referred to in this presentation are presented as illustrations of the manner in which response time would be enhanced when queries qualify to offload to the accelerator. Actual query workloads may offload at substantially different rates depending on the degree to which the queries meet first release requirements for query block offload.
 - This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
 - All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Workload A Query Performance

Query	ET sec. on z	CPU sec. on z	ET secs with Accel. Avail	CPU secs with Accel. Available	% Reduction in ET	Improvement Factor
WorkloadA-Q01	317.35	1149.61	4.35	0.02	98.63%	72.95
WorkloadA-Q02	11.41	2.06	6.79	2.07	40.49%	1.68
WorkloadA-Q03	12.77	4.61	8.33	3.42	34.77%	1.53
WorkloadA-Q04	57.45	169.98	1.16	0.02	97.98%	49.53
WorkloadA-Q05	683.33	99.75	13.69	0.02	98.00%	49.91
WorkloadA-Q06	27.26	6.83	0.83	0.04	96.96%	32.84
WorkloadA-Q07	55.44	54.46	2.93	0.02	94.72%	18.92
WorkloadA-Q08	1.84	0.10	1.75	0.11	4.89%	1.05
WorkloadA-Q09	17.14	1.02	0.4	0.02	97.67%	42.85
WorkloadA-Q10	95.87	217.07	13.27	0.02	86.16%	7.22
WorkloadA-Q11	2.21	2.34	0.18	0.03	91.86%	12.28
WorkloadA-Q12	0.28	0.03	0.14	0.02	50.00%	2.00
WorkloadA-Q13	41.25	1.74	1.12	0.02	97.28%	36.83
WorkloadA-Q14	2.21	0.14	1.09	0.03	50.68%	2.03
WorkloadA-Q15	439.51	28.98	1.29	0.03	99.71%	340.71
WorkloadA-Q16	303.06	5.70	1.67	0.02	99.45%	181.47
WorkloadA-Q17	305.47	96.31	5.17	0.02	98.31%	59.09
WorkloadA-Q18	630.51	2339.65	420.18	1116.94	33.36%	1.50
WorkloadA-Q19	460.53	6.45	2.63	0.03	99.43%	175.11
WorkloadA-Q20	2427.46	6392.94	8.42	0.03	99.65%	288.30
WorkloadA-Q21	1882.22	791.77	1.72	0.03	99.91%	1094.31
WorkloadA-Q22	29.38	110.56	1.54	0.02	94.76%	19.08
WorkloadA-Q23	730.28	698.70	0.47	0.02	99.94%	1553.79
WorkloadA-Q24	2351.87	2231.17	1.32	0.04	99.94%	1781.72
WorkloadA-Q25	1645.72	704.58	1.65	0.03	99.90%	997.41
WorkloadA-Q26	116.36	9.19	2.8	0.04	97.59%	41.56

Workload B Query Performance

Query	ET sec. on z	CPU sec. on z	ET secs. With Accel available	CPU secs. With Accel available	% Reduction in ET	Improvement Factor
WorkloadB_Q01	40.98	143.39	1.18	0.02	97.12%	34.73
WorkloadB_Q02	480.73	843.38	3.64	0.03	99.24%	132.07
WorkloadB_Q03	43.10	158.26	2.93	0.02	93.20%	14.71
WorkloadB_Q04	690.82	1312.67	2.38	0.02	99.66%	290.26
WorkloadB_Q05	1090.35	4.18	0.75	0.02	99.93%	1453.80
WorkloadB_Q06	449.28	757.59	2.46	0.02	99.45%	182.63
WorkloadB_Q07	328.49	784.10	2.47	0.02	99.25%	132.99
WorkloadB_Q08	357.94	824.28	2.63	0.02	99.27%	136.10
WorkloadB_Q09	474.47	1286.73	3.95	0.02	99.17%	120.12
WorkloadB_Q10	640.27	1858.90	3.89	0.02	99.39%	164.59
WorkloadB_Q11	883.42	65.09	2.29	0.02	99.74%	385.77
WorkloadB_Q12	573.04	1300.95	5.26	0.02	99.08%	108.94
WorkloadB_Q13	272.74	327.85	0.78	0.02	99.71%	349.67

DB2 Data Sharing

- Testing took place using 1 DB2 data sharing group
- Two z10 LPARs on a single CEC each contained one DB2 data sharing member, DSNA & DSNB
- The Smart Analytics Optimizer was installed on DSNA on SYSA, but was started on both data sharing members
- To test access from both data sharing members to the accelerator, multiple jobs each containing multiple SQL steps were submitted on LPARs SYSA and SYSB

DB2 Data Sharing...

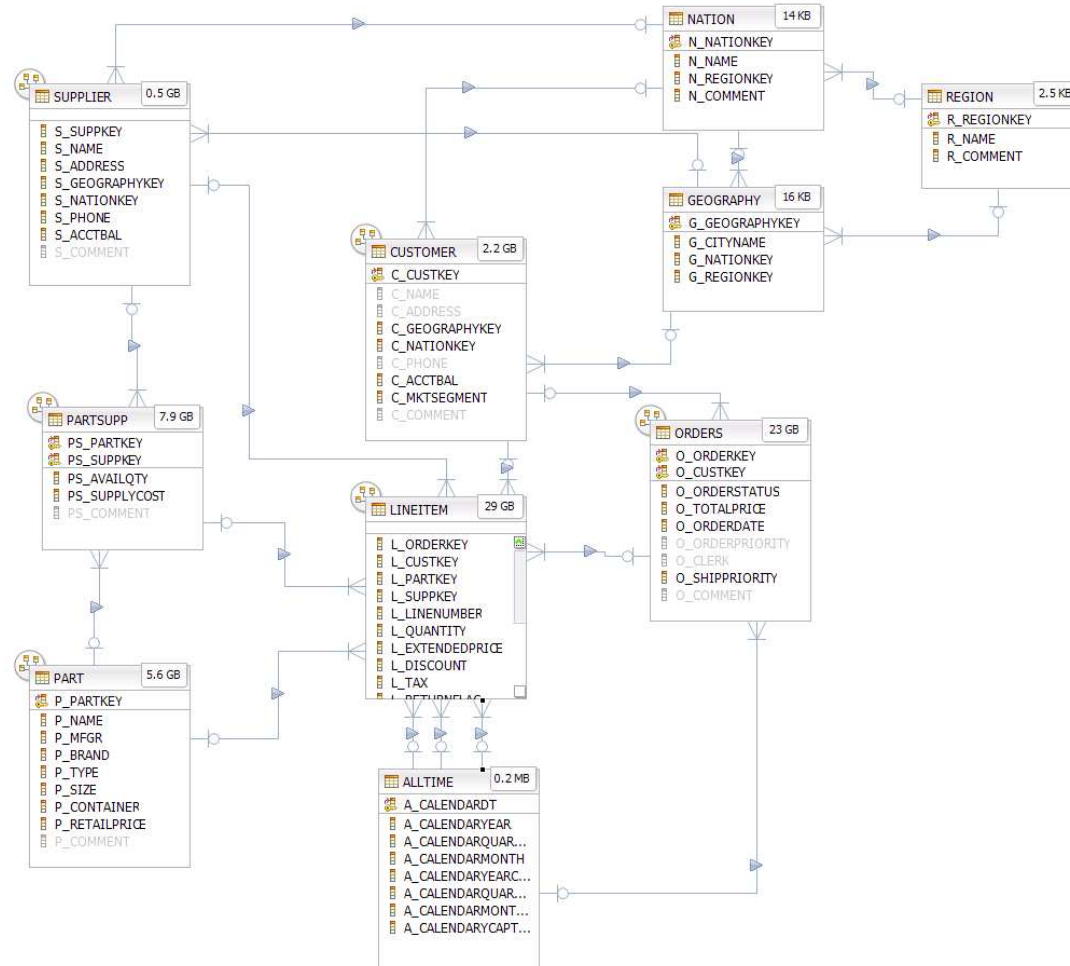
- Queries from both DSNA and DSNB were offloaded to the accelerator and interspersed in the accelerator queue
- Since the accelerator processes only one Query block at a time, the other Queries from the multiple jobs queued on the accelerator waiting to execute
- When the queue length exceeded a threshold, additional queries arriving at the accelerator received an internal error and then executed on z
- Subsequent job steps after executing on z ran on the accelerator as long as the queue length did exceed the threshold

Alternative Data Mart Definition for Workload A

- At the lab's suggestion, I overrode the default dimension table definition to fact table for all of the tables that had caused me to create multiple data models
- The Smart Analytics Optimizer Studio data mart diagram is shown on the next page
- Estimated compressed storage for this single data model is no greater than the Lineltem data mart I was using previously
- All Workload A queries except the shortest running queries offloaded to the accelerator and the performance of these queries is shown on page 46

Workload A Single Data Mart definition with 6 Fact tables

Estimated size: 67 GB



Workload A Query Performance with a single Data Mart and multiple Fact tables

Query	ET sec. on z	CPU sec. on z	ET secs with Accel. Avail	CPU secs with Accel. Available	% Reduction in ET	Improvement Factor
WorkloadA-Q01	317.35	1149.61	3.750	0.030	98.82%	84.63
WorkloadA-Q02	11.41	2.06	7.030	9.700	38.39%	1.62
WorkloadA-Q03	12.77	4.61	9.590	16.110	24.90%	1.33
WorkloadA-Q04	57.45	169.98	0.960	0.020	98.33%	59.84
WorkloadA-Q05	683.33	99.75	2.660	0.020	99.61%	256.89
WorkloadA-Q06	27.26	6.83	0.430	0.020	98.42%	63.40
WorkloadA-Q07	55.44	54.46	0.790	0.020	98.58%	70.18
WorkloadA-Q08	1.84	0.10	0.170	0.020	90.76%	10.82
WorkloadA-Q09	17.14	1.02	0.360	0.020	97.90%	47.61
WorkloadA-Q10	95.87	217.07	2.750	0.020	97.13%	34.86
WorkloadA-Q11	2.21	2.34	0.170	0.020	92.31%	13.00
WorkloadA-Q12	0.28	0.03	0.170	0.020	39.29%	1.65
WorkloadA-Q13	41.25	1.74	1.080	0.030	97.38%	38.19
WorkloadA-Q14	2.21	0.14	1.050	0.030	52.49%	2.10
WorkloadA-Q15	439.51	28.98	8.970	7.790	97.96%	49.00
WorkloadA-Q16	303.06	5.70	1.230	0.030	99.59%	246.39
WorkloadA-Q17	305.47	96.31	1.890	0.030	99.38%	161.62
WorkloadA-Q18	630.51	2339.65	31.000	0.030	95.08%	20.34
WorkloadA-Q19	460.53	6.45	1.510	0.030	99.67%	304.99
WorkloadA-Q20	2427.46	6392.94	2.450	0.720	99.90%	990.80
WorkloadA-Q21	1882.22	791.77	1.230	0.030	99.93%	1530.26
WorkloadA-Q22	29.38	110.56	0.590	0.020	97.99%	49.80
WorkloadA-Q23	730.28	698.70	0.400	0.020	99.95%	1825.70
WorkloadA-Q24	2351.87	2231.17	1.200	0.030	99.95%	1959.89
WorkloadA-Q25	1645.72	704.58	1.570	0.030	99.90%	1048.23
WorkloadA-Q26	116.36	9.19	0.830	0.030	99.29%	140.19

Q&A

