• IBM DB2 Analytics Accelerator for z/OS V2.1 is building upon the System z platform strength and extends the value of the OLTP data for the customer by enabling highly efficient and performing reports and analytics on this data.
- It enables decision makers to perform business analysis they had never considered in the past to analyze trends, predict outcomes, and produce better business results.

Key competitive strength are:

Performance: Unprecedented response times to enable 'train of thought' analysis that were frequently blocked before by poor query performance.

Extend System z QoS beyond OLTP data: By hosting data on System z the traditional QoS values such as security, integrity, reliability and availability are applied to data warehouses as well.

Reduced TCO for hosting Data Warehouses: Combined effect of consolidating DW/DM workload onto a single platform (reuse of skills, tools, operating procedures) and executing very resource intensive queries on a price/performance optimized platform.

NOTE - TeraData installations are the prime target for applying this competitive strength.

Integration: Connecting to DB2 through deep integration is providing transparency to all apps. Implemented as a formal, internal DB2 component this is resulting in consolidated database administration and transparency to database applications, despite heterogeneous platform nature.

Self-managed workloads: Queries are executed in the most efficient way.

Transparency: Applications connected to DB2 are entirely unaware of the Accelerator. This enables the use of existing applications without any change as well as developing new applications that connect to DB2 as the one database management system for all.

IBM DB2 Analytics Accelerator for z/OS will allow the customer:
- To extend the use of operational platform data to perform business analysis and daily reporting.
- To run fast and cost effective business reporting on a single platform (e.g. by integrating OLTP data into business analytics and reporting).
- To substantially reduced operational costs through - removing the need for complex query tuning offloading complex and long running queries to a specialized subsystem. to make room for additional workloads due to off-loading capabilities.
- To simplify administration: appliance hands-free operations, eliminating few database tuning tasks.

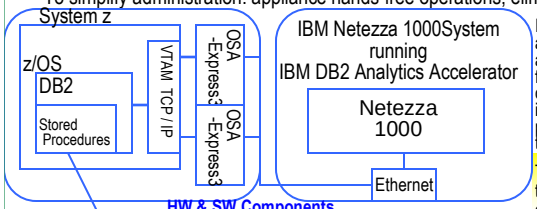

HW & SW Components

IDAA **stored procedures** are the administration interface for your accelerators. When you invoke a function from IDAA Studio, the corresponding stored procedure is called. The stored procedures provide functions that are related to tables and accelerators.

DB2 Analytics Accelerator Studio

DB2 isolation levels are ignored by IBM DB2 Analytics Accelerator for z/OS.

To make DB2 for z/OS route queries to an accelerator, you must enable query redirection using the CURRENT QUERY ACCELERATION special register. Alternatively, you can set the QUERY_ACCELERATION ZPARM in DB2 for z/OS.

• A single accelerator can be shared by multiple DB2 subsystems.
NOTE: A single DB2 for z/OS subsystem can also be connected to more than one accelerator.

IBM DB2 Analytics Accelerator for z/OS supports the following subsystem configurations:
- Multiple subsystems, each of which in a separate logical partition (LPAR)
- Multiple subsystems in a common LPAR
- Multiple subsystems that make up a data sharing group (subsystems in different LPARs, on different Central Processing Complexes (CPCs))

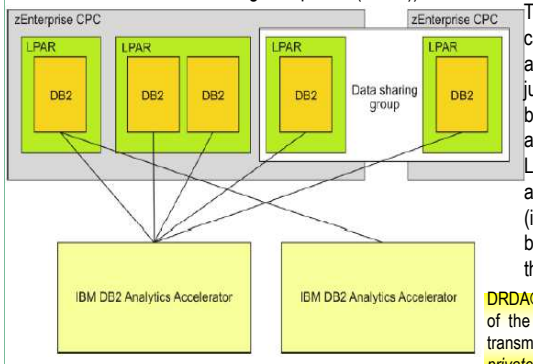

Possible connections

This shows that DB2 subsystems can *share* a single accelerator as well as connect to more than just one accelerator. The leftmost box in the figure, which represents a single subsystem in a separate LPAR, is connected to two accelerators. All DB2 subsystems (including the one in the leftmost box) share the one accelerator on the left.

DRDA® is used for the transmission of the query. The proprietary DB2 transmission protocol known as the *private protocol* is **not** supported.

Cheat Sheet #85 zTidBits

z/OS IBM DB2 Analytics Accelerator

IBM DB2 Analytics Accelerator FMID: HAQT210

Minimum Hardware

IBM System zEnterprise™ 196 (2817)
or
IBM System zEnterprise 114 (2818)

IBM Netezza® 1000-3 with NA 5725-E46
or
IBM Netezza 1000-3 with EMEA 5725-E47
or
IBM Netezza 1000-6 5725-E48
or
IBM Netezza 1000-12 5725-E49

Workstation (Intel® or compatible) for IBM DB2 Analytics Accelerator Studio. This workstation must have connectivity to the zEnterprise server.

Based on its own set of heuristics, DB2 classifies an acceleration of the query as favorable. That is, DB2 "expects" a much shorter query response time by routing the query to an accelerator.

IBM DB2 Analytics Accelerator Studio is the graphical administration interface for the product and is delivered on the product DVD. It consists of a set of Eclipse plug-ins to be added to IBM Data Studio.

Instead of installing IBM DB2 Analytics Accelerator Studio from the product DVD, you can also add the plug-ins to an existing IBM Data Studio 2.2.1.0 or 2.2.1.1 installation. Both IBM Data Studio products are free of charge.

To assist you in determining if you have the **recommended service** for IBM DB2 Analytics Accelerator for z/OS, V2.1 installed on your system, you can use the SMP/E REPORT MISSINGFIX command in conjunction with FIXCAT HOLD DATA;

See: #17 zNibbler (SMP/E's FIXCAT HOLDDATA)

When you select tables for query acceleration, information about these tables is added to the IBM DB2 Analytics Accelerator for z/OS system tables in the DB2 for z/OS catalog of the connected subsystem. Each entry (system table row) serves as a link between DB2 and the table on the accelerator. An entry contains, among other information, the table name.

Important: Queries that are routed to an accelerator are **not cached** in the DB2 Dynamic Statement Cache.

Connecting to a data server: IBM DB2 Analytics Accelerator for z/OS needs access to your data servers to read the database catalogs and invoke the DB2 stored procedures for the accelerators.

Creating a database connection profile: IBM DB2 Analytics Accelerator Studio - the user interface for administering Accelerators and it connects to your DB2 for z/OS data server using Java Database Connectivity (JDBC), like other database clients.
•The configuration parameters for accessing a DB2 for z/OS data server are stored in database connection profiles.
• Create a database connection profile for each DB2 for z/OS data server with an accelerator.
• When a connection to a database has been established through one of your database connection profiles, you can see all currently deployed accelerators in the Object List Editor.
-You access these interface elements from the Accelerator Perspective in IBM DB2 Analytics Accelerator Studio.

Defining the data to load into an accelerator: Defining the data to load into an accelerator basically means selecting the proper tables and choosing favorable distribution keys and organizing keys.

Checking the table size: When you select a table for an accelerator in the Add Tables wizard, IBM DB2 Analytics Accelerator Studio displays the table size.
• This table size is calculated on the basis of the values in the CARDF and AVGROWLEN columns of the SYSIBM.SYSTABLES table.
- The metrics come from DB2 for z/OS.
- Comparing this size with the available disk space on your accelerator, you can roughly assess how much space the table will occupy on your accelerator.
NOTE: DB2 for z/OS and the accelerator use completely different compression algorithms, so the size of a table in DB2 might differ from the size of the same table on the accelerator.

NOTE: Plans for query processing must use packages rather than database request modules (DBRMs)

Loading data into selected tables: To enable users to run accelerated queries against selected database tables, you must load the table data into the empty tables whose layout or definition has been copied to an accelerator therefore, you create a snapshot copy of your existing DB2 for z/OS data on the accelerator.
• The optimizer of your database management system calculates the expected response times for incoming queries.
- If an accelerator can process a query against the selected tables faster than the database management system, the query is routed to the accelerator and evaluated against the populated tables.
NOTE: Successful queries against tables on an accelerator are possible only if the tables contain data. Therefore, you load the tables after their definition (empty structure) has been copied to the accelerator.

Enabling tables for query acceleration: Query acceleration can be switched on & off for each table on an accelerator allowing you to do maintenance work on a particular table while query acceleration remain active for other tables on the same accelerator.

Disabling query acceleration for a table: Disabling a table, you can prevent accelerated queries against a table where it might be necessary when a query must be run against very recent data, which has not yet been copied to the accelerator.
NOTE: A table can be added to one accelerator only.
• A table on an accelerator represents a table in your database and by adding a table in IBM DB2 Analytics Accelerator Studio, you instruct IBM DB2 Analytics Accelerator for z/OS to copy the definition of that table (an empty table) to an accelerator.
• Queries including this table are then routed to the accelerator provided that the requirements for query redirection are met.
- A query can only be routed to an accelerator if the query refers to a subset or all of the tables on the accelerator.
• A query can only be accelerated if all tables that are referenced by the query reside on the same accelerator, if the tables on the accelerator are loaded, & if query acceleration is enabled for these tables.
- A query can reference tables indirectly if the table references a view or a table alias.
- if the query references a view that joins two tables, both tables must be present on the accelerator.
• When tables are added to an accelerator, DB2 for z/OS keeps track of these tables in its system catalog.
• The system table entries enable the optimizer of the database management system to complete the following tasks:
- Query matching, that is, comparing an incoming query with the information stored in the system tables to check whether the query can be handled by IBM DB2 Analytics Accelerator for z/OS.
- Query evaluation, that is, calculating the estimated query response time; if a performance gain can be predicted, the query is sent to IBM DB2 Analytics Accelerator for z/OS; if not, the query is handled by the database management system itself.

Queries can be redirected to an accelerator only if certain conditions are met.
NOTE: There are also adverse conditions that prohibit query acceleration.
A query can be redirected to IBM DB2 Analytics Accelerator for z/OS if the following general conditions are met:
1. The accelerator is in an operational mode and has been started.
2. DB2 for z/OS has been configured so that the query will be routed to a connected accelerator provided that this query meets all other conditions.
3. IBM DB2 Analytics Accelerator for z/OS supports the SQL expressions that are used in the query.
4. All data that is referenced by the SQL code in the query, such as the contents of tables and views, is available on one and the same accelerator.

NOTE: To satisfy this condition, the accelerator must have been loaded with the data in question.