



Welcome to the Virtual IMS user group newsletter. The Virtual IMS user group at www.fundi.com/virtualims is an independently-operated vendor-neutral site run by and for the IMS user community.

Virtual IMS user group presentation

The latest webinar from the Virtual IMS user group was entitled, “Leveraging IBM DB2 Analytics Accelerator to Analyze your IMS Data”, and was presented by Rocket Software’s Wayne Morton.

Wayne is an experienced mainframe systems architect having spent more than 20 years working with zEnterprise software. Over the years, he has garnered experience in DB2, IMS, and CICS. Morton was part of the Research and Development division at BMC Software, and most recently has managed technical field operations for NEON Systems and Progress Software. In his role with Rocket Software, he is focused on customer engagement to help identify

Transactional Processing

Analytics Workload



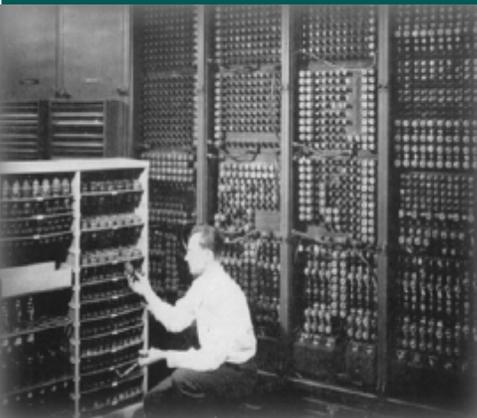
Figure 1: DB2 Analytics Accelerator and DB2 for z/OS

cost-efficient methods for accessing mainframe data.

Wayne started the presentation by saying IDAA stands for IBM DB2 Analytics Accelerator. He went on to say that DB2 Analytics Accelerator and DB2 for z/OS make up a self-managing, hybrid workload-optimized database management system that

Contents:

Virtual IMS user group presentation	1
Meeting dates	4
IMS articles	4
IMS news	4
About the Virtual IMS user group	4



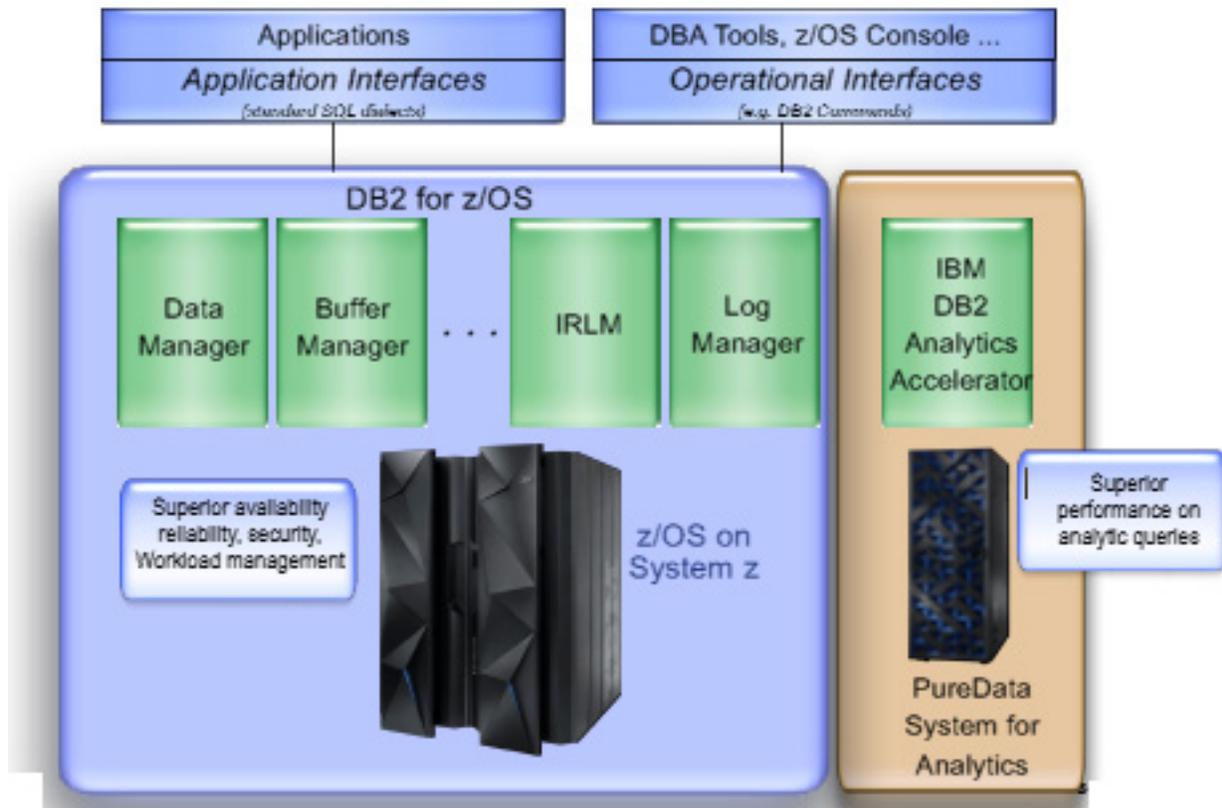


Figure 2: Deep integration with zEnterprise

runs query workloads in the most efficient way, so that queries are executed in the optimal environment for greatest performance and cost efficiency.

The hybrid computing platform on zEnterprise:

- Supports transaction processing and analytics workloads concurrently, efficiently, and cost-effectively
- Delivers industry leading performance for mixed workloads.

As Figures 1 and 2 clearly show, the IDAA doesn't use any expensive mainframe

MIPS because it runs on PureData for Analytics (powered by Netezza technology).

The advantages of this set up are: unprecedented response times for "right-time" analysis; complex queries in seconds rather than hours; transparent to the application; inherits all System z DB2 attributes; no need to create or maintain indices; eliminate query tuning; fast deployment and time-to-value.

The tables need to be defined and deployed to IDAA before data is loaded and queries sent to it for

processing. Definition means identifying tables for which queries need to be accelerated. Deployment is making tables known to DB2, ie storing table metadata in the DB2 and PureData catalog. IDAA Studio guides you through the process of defining and deploying tables, as well as invoking other administrative tasks. IDAA Stored Procedures implement and execute various administrative operations such as table deployment, load and update, and serve as the primary administrative interface to IDAA from the outside world including IDAA Studio.

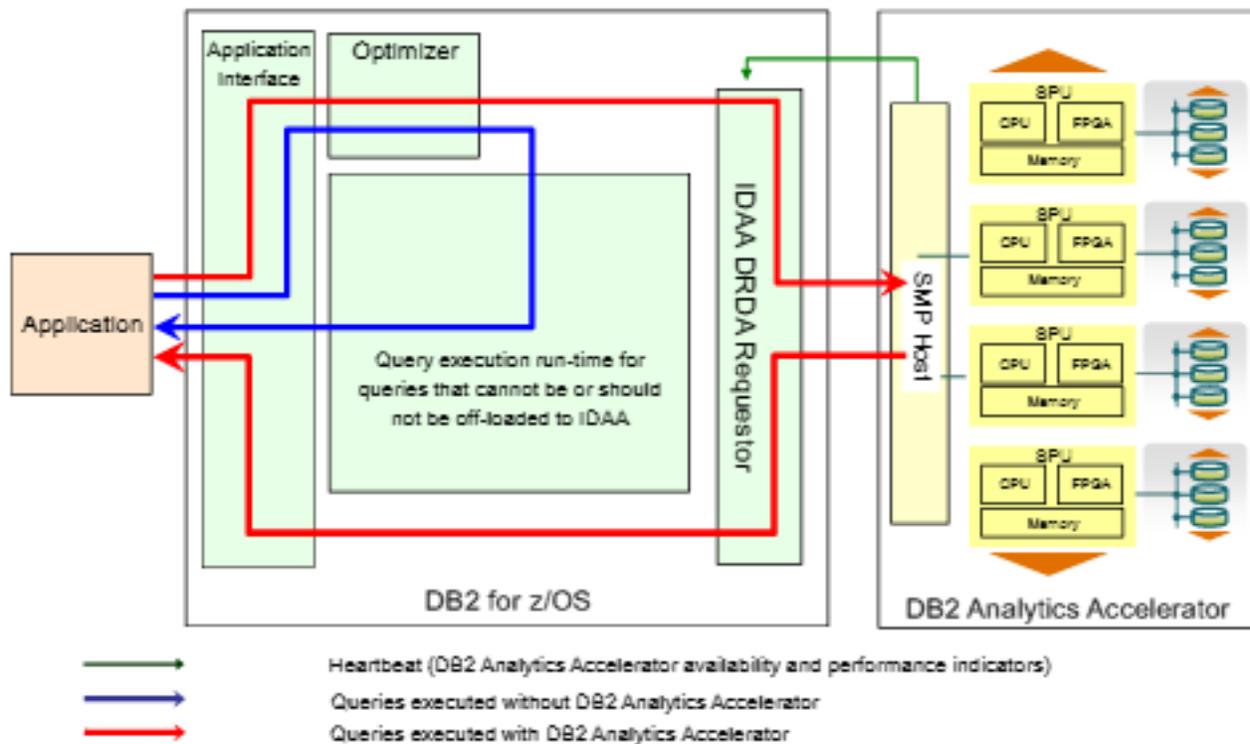


Figure 3: Query execution process flow

Query execution process flow is illustrated in Figure 3. The blue line shows what happens to queries executed without DB Analytics Accelerator. The red line shows queries executed with DB2 Analytics Accelerator.

But the tool isn't just for DB2. Customers want to learn more about their investment in the Accelerator and maximize its use in their environment. These customers are looking at creative ways to exploit the Accelerator for IMS, VSAM, SMF data, and non-z/OS data.

There are three different areas where tools can provide value:

- Assessment – do I have a workload that would benefit from the Accelerator?
- Optimization – can I optimize the workload to take advantage of the Accelerator?
- Administration – can I manage the Accelerator more effectively?

Wayne went on to explain how data gets loaded into the Accelerator. There's the standard DB2 UNLOAD utility that extracts data. It places tables in read-only mode until the Accelerator load is done (when using the option for transactional consistent data). The Accelerator

reflects DB2 data at the point in time the load was initiated. You could, optionally, use CDC (Change Data Capture) for near-real-time replication to the Accelerator

The new IBM DB2 Accelerator Loader V1.1 provides group consistent load and external 'dual' load. It's built for performance and usability. The Loader populates data in the Accelerator. It can load DB2 and non-DB2 data. It loads in parallel to avoid application downtime. And it can load to an historical point-in-time.

The Accelerator Loader can load data from a file in one of two ways – dual external load and Accelerator only.

With dual external load, it loads data into both DB2 and the Accelerator in parallel. With Accelerator only, the Accelerator Loader loads directly into the Accelerator (no load in DB2).

The user is responsible for building the load file. Extracted data can come from various sources, eg IMS, VSAM, Oracle, etc. The file must be compatible for input into the DB2 LOAD utility. Field specification must describe the input data format. This must be compatible with the DB2 LOAD utility. Null-able columns or not null with default can be missing from input data. Defaults will be supplied by Accelerator Loader

A copy of Wayne's presentation is available for download from the Virtual IMS user group Web site at www.fundi.com/virtualims/presentations/IDAALoaderDec14.pdf.

You can see and hear the whole user group meeting by downloading the WMV file from www.fundi.com/virtualims/presentations/2014-12-09meeting.wmv.

Meeting dates

- On 10 February 2015, we have GT Software's Dusty Rivers.

- On 14 April 2015, it's a speaker from DataVantage.

Recent IMS articles

Accessing IMS Data Through the JDBC API by Richard Tran in *IBM Systems Magazine* (October 2014). You can find the article at <http://www.ibmssystemsmag.com/mainframe/administrator/ims/JDBC-API/>.

IMS news

Compuware has announced Version 12.4 of Abend-AID, its fault management tool. The new release features tighter cross-product integration, modernized utilities, and usability enhancements. There's deeper integration between Abend-AID and File-AID within the Eclipse plug-ins, which facilitates faster problem solving. Abend-AID Fault Analytics now shares common Web services with iStrobe to help users simplify and automate common tasks. Compuware Find & Fix functionality has been expanded to include IMS, Batch, and DB2 components along with existing CICS support.

About the Virtual IMS user group

The Virtual IMS user group was established as a way for individuals using IBM's IMS hierarchical database and transaction processing systems to exchange information, learn new techniques, and advance their skills with the product.

The Web site at www.fundi.com/virtualims provides a central point for coordinating periodic meetings (which contain technically-oriented topics presented in a webinar format), and provides articles, discussions, links, and other resources of interest to IBM IMS practitioners. Anyone with an interest in IMS is welcome to join the Virtual IMS user group and share in the knowledge exchange.

To share ideas, and for further information, contact trevor@itech-ed.com.

The Virtual IMS user group is free to its members.