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, "Replicating IMS Data to IDAA and PureData Analytics", and was presented by Scott Quillicy, the CEO and founder of SQData.

SQData is an Addison, TX-based software company that specializes in high-performance data movement and changed data capture (CDC) for IMS, VSAM, and relational databases. He has over 30 years of database experience and is considered an expert in database replication strategy and deployment. Scott has managed a number of large data integration projects, many of which involved the near-real-time data capture of IMS databases to a variety

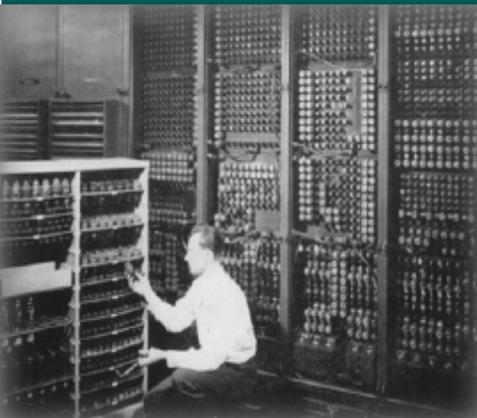**Figure 1: DB2 Analytics Accelerator and DB2 for z/OS**

of downstream databases/applications.

Scott started his presentation by looking at reasons to replicate IMS to IDAA (IBM DB2 Analytics Accelerator) and PDA (PureData Analytics). They were: to provide a method of analyzing data outside of IMS; to provide real-time business intelligence

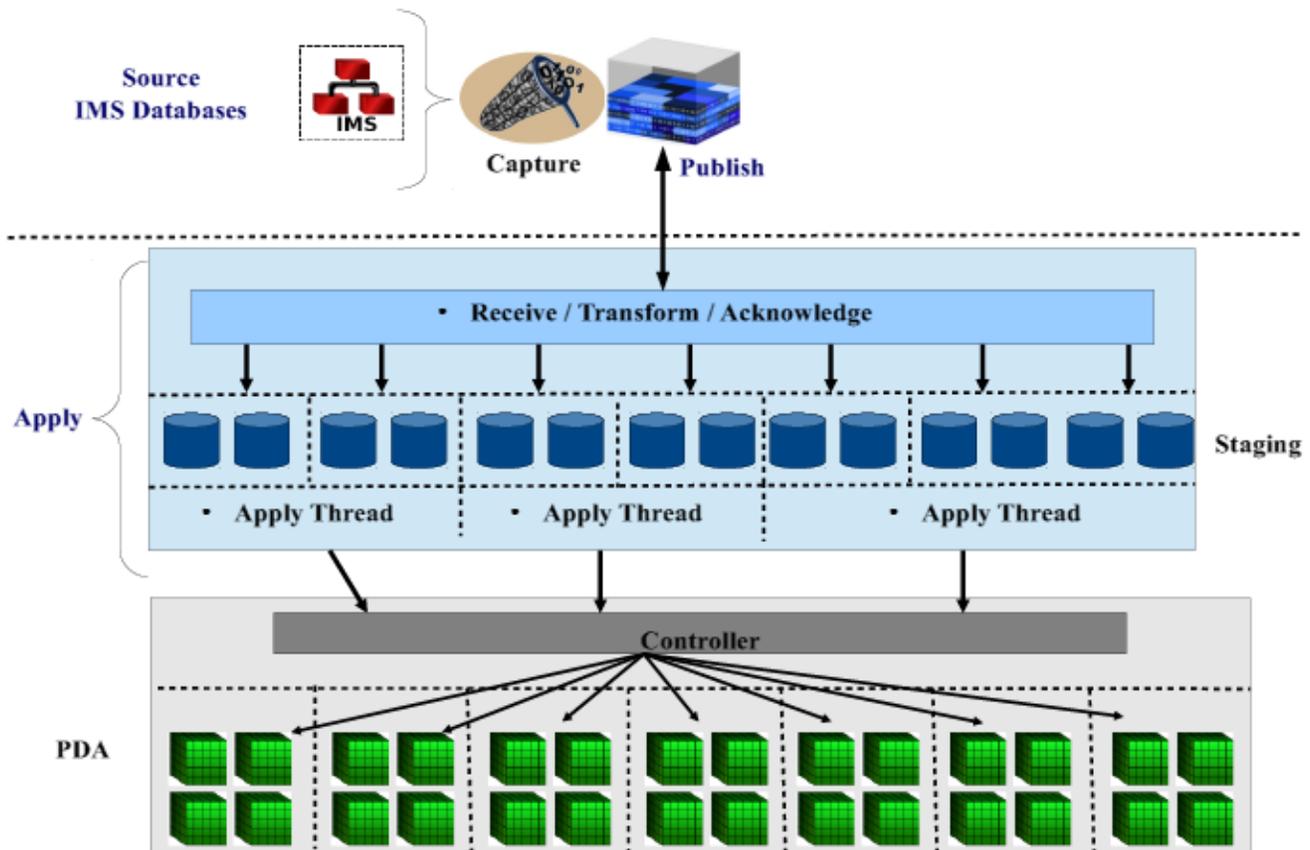/ advanced analytics; because bulk loads can

Figure 2: IMS to PDA replication

be resource intensive; to combine with data from other applications; to save significant CPU cycles for intense queries; and to compliment established Data Warehouse(s).

PureData Analytics uses a Netezza Appliance. IBM DB2 Analytics Accelerator uses a Netezza Appliance coupled to DB2 (see Figure 1).

Scott went on to describe the role of ETL and CDC ETL (Extract, Transform, Load), as a full data extract / load. The data transformation logic is defined in this step. It's an iterative process that

must be fast and efficient. And it should minimize data landing. On the other hand, CDC (Changed Data Capture) moves only the data that has changed, can re-use data transformation logic from ETL, and is near-real-time / has deferred latency.

When it comes to performing the initial load, Scott suggested:

- Transformation / mapping logic is done here and reused in CDC.

- Should be able to run against live sources.

- Make sure to truncate before loading to avoid duplicates.

- May be used in lieu of incremental updates (CDC).

- NZLOAD is the native loader and allows for rapid loading of PDA tables.

- IDAA Loader (recommended component) is an IBM product offering that allows simultaneous loading into DB2 and IDAA, and also allows IDAA-only and AOT loads.
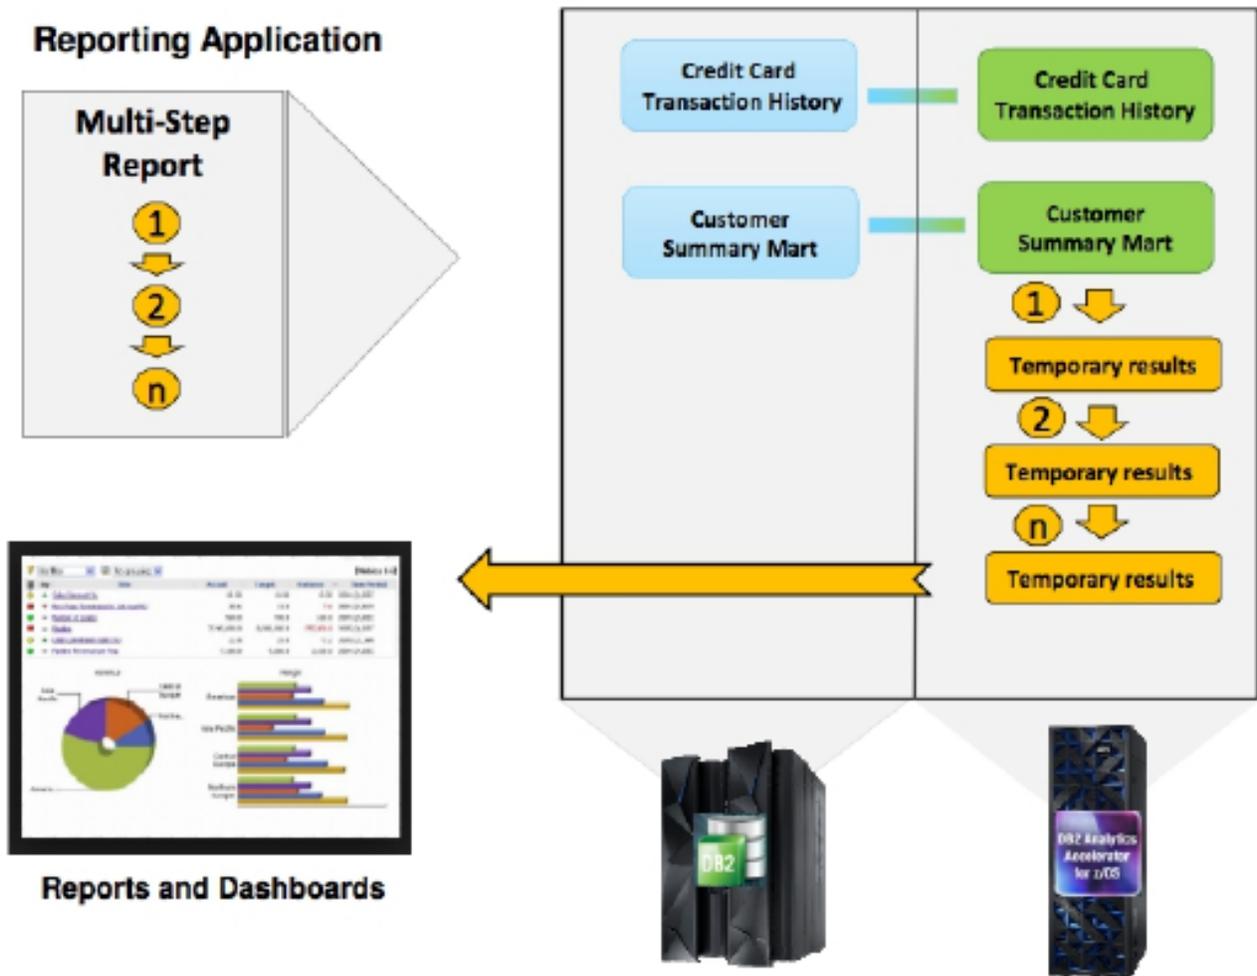
Figure 3: Query workload after AOTs

When it comes to Insert, Update, and Delete behaviour:

- Updates – Delete / Insert Pairs.

- Inserts append data very quickly to the end of the file. Speed is based on the number of rows being inserted.

- Deletes: must scan the entire file; select row – update with delete flag; replace row in place; the groom process cleans up files; and speed is based on the size of the file and the number of rows being deleted.

Scott's replication strategy:

- General approach, apply in 'batches' rather than continuously.

- Your expectations on latency should be minutes rather than sub-second.

- Time-series data tracks the lineage of changes and is commonly deployment to track customer tendencies, campaign effectiveness, correlation of events. All changes are applied as inserts, and the options are batch or continuous feed.

- Synchronized data – the source and target match at any given time, and inserts, updates and deletes must be
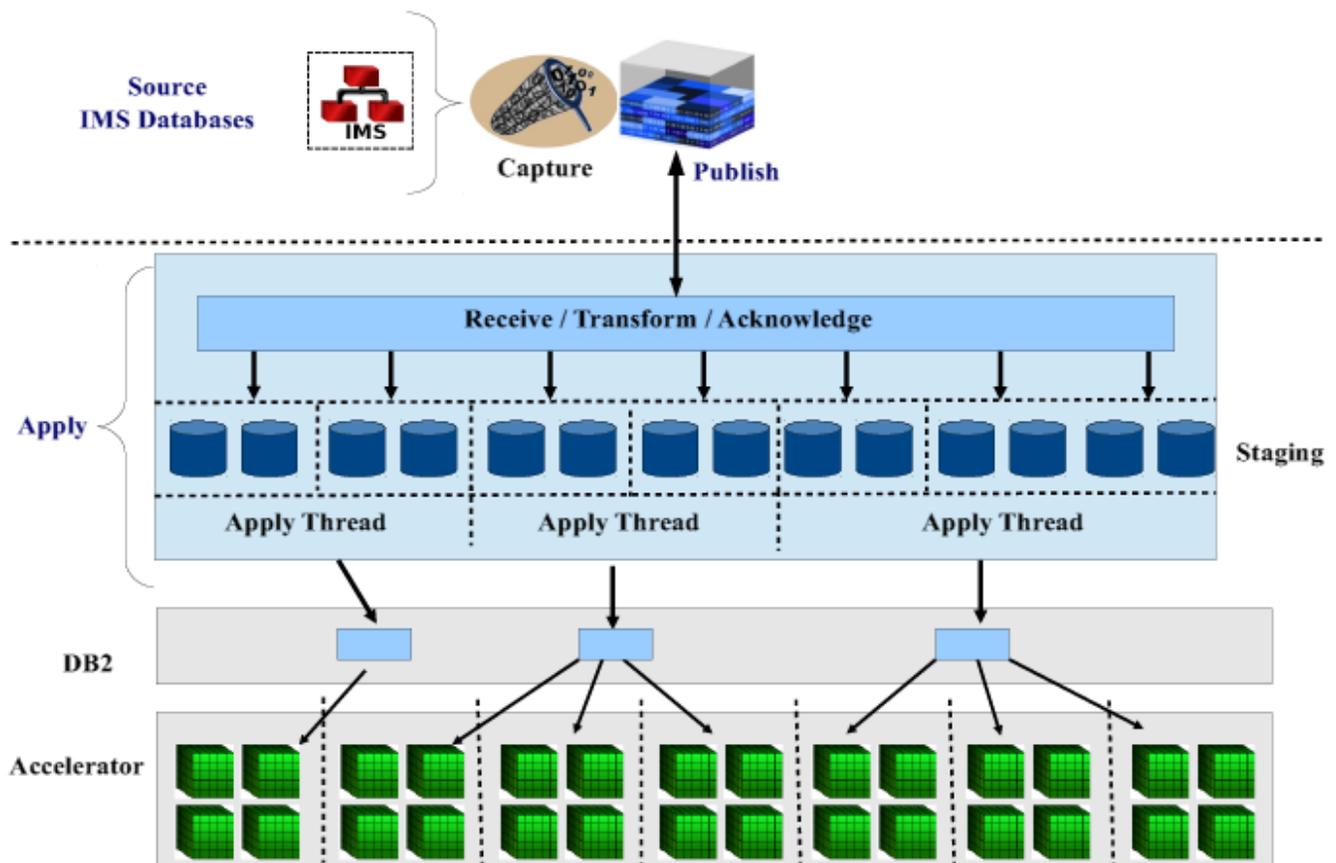
Figure 4: IMS to DB2AA replication

processed. The only option is batch.

When it comes to target key selection, the golden rule is that good distribution equals good performance. The objective are: all tables are distributed across all active database blades; all queries run in parallel against all active database blades; and all loads run parallel against all active database blades.

An example of IMS to PDA replication is shown in Figure 2.

For IMS to DB2AA replication, the Accelerator

must know about apply processes. Accelerator Only Tables (AOTs) allows Update DML against Tables in the Accelerator; the apply process can perform Inserts/Deletes through DB2; and although throughput is currently fine, it will get better in the future. A query workload using AOTs is illustrated in Figure 3.

An example of IMS to DB2AA replication is illustrated in Figure 4.

A copy of Scott Quillicy's presentation is available for download from the Virtual IMS user group Web

site at www.fundi.com/virtualims/presentations/IMSReplicationAug15.pdf.

You can see and hear the whole user group meeting by downloading the WMV file from www.fundi.com/virtualims/presentations/2015-08-11meeting.wmv.

Meeting dates

- On 6 October 2015, we have BMC's Glenn Witt.

- The final meeting of this year is on 1 December 2015.

IMS news

HostBridge Technology has announced HostBridge Redis for z/OS. The product enables customers to put their mainframe data to work in analytics, and mobile and cloud architectures. Redis is the in-memory NoSQL data store used by cloud and mobile apps. HostBridge Redis for z/OS is a modernizing extension for transaction processing systems such as IMS and CICS, and database systems such as DB2 and VSAM. Using HostBridge Redis, customers can share their data with applications on any platform, inside or outside the enterprise. More information can be found at https://www.hostbridge.com/index.php/products/redis-for-zos.

Recent IMS articles

IMS Security Frequently Asked Questions by Maida Snapper in *IBM Systems Magazine* (August 2015). You can find the article at http://www.ibmsystemsmag.com/mainframe/tipstechniques/miscellaneous/IMS-Security-FAQ/.

DIY: Migrating VSAM files to IMS Database by Carson Tsai, Rita Shih, and Jack Yuan in *IBM Systems Magazine* (July 2015). You can find the article at http://www.ibmsystemsmag.com/mainframe/administrator/ims/DIY-migration/.

Connector Architecture by Richard Tran in *IBM Systems Magazine* (June 2015). You can find the article at http://www.ibmsystemsmag.com/mainframe/administrator/ims/Java-EE-Connector/.

Like us on Facebook

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.