



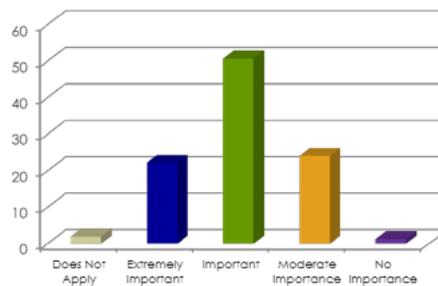
Welcome to the Virtual IMS user group newsletter. The Virtual IMS user group at [www.fundi.com/virtualims](http://www.fundi.com/virtualims) is an independently-operated vendor-neutral site run by and for the IMS user community.

## Virtual IMS user group presentation

The latest webinar from the Virtual IMS user group was entitled, "Putting IMS/DB in Business Analytics/Intelligence", 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. Early in his career, he worked for Wal~mart, first as a DBA, and later as a systems programmer. Over the years, he has garnered experience in DB2, IMS DB, 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

- How important is it for your organization to integrate non-relational data with a business analytics or business intelligence initiative?



Source: Enterprise Systems Media: "Mainframe Data and Business Analytics" - D. Yost, November 2012



more than 70% signified high level of importance for integrating non-relational data with analytics

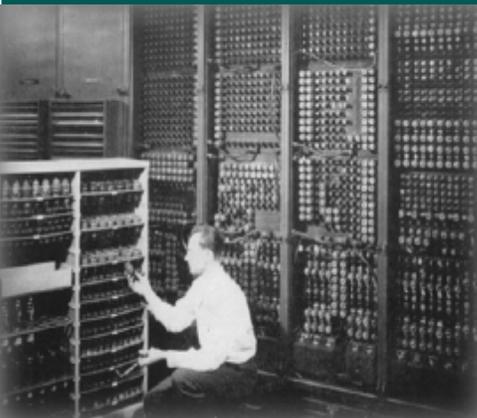
**Figure 1: Analytics and mainframe data**

Progress Software. In his role with Rocket Software, he is focused on customer engagement to help identify cost-efficient methods for accessing mainframe data.

Wayne Morton started by saying that we have seen phenomenal growth associated with business analytics and business intelligence applications.

### Contents:

Virtual IMS user group presentation	1
Meeting dates	5
IMS news	5
IMS articles	5
About the Virtual IMS user group	5



According to IDC Research, the BI and analytic tools segment grew 13.2% in 2011, and will continue to grow at a 9.8% compound annual growth rate (CAGR) through 2016 to reach \$50.7 billion.

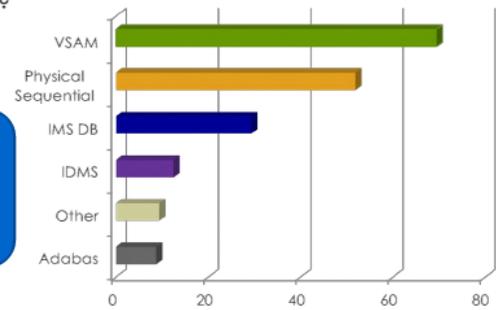
To quote IDC’s Dan Vesset: “Driven by the attention-grabbing headlines for Big Data and more than three decades of evolutionary and revolutionary developments in technology and best practices, the business analytics software market has crossed the chasm into the mainstream mass market.”

This is particularly relevant when it comes to mainframe organizations. Figure 1 shows that more than 70% of the organizations represented in this survey indicate that integrating mainframe non-relational data with business analytics or business intelligence is either extremely important or important.

But where does the data reside, and to what extent? Figure 2 answers those questions. As you can see there is still a lot of important data stored in VSAM, sequential or flat files, IMS, IDMS, Adabas, and other similar databases. In fact, 351 respondents indicated the data they need exposed

- Which non-relational data source will need to be exposed to your business analytics or business intelligence application?

69% of respondents listed VSAM as the most prevalent non-relational data source needed for analytics



**Figure 2: Analytics and non-relational data**

to business analytics or business intelligence applications resides in VSAM files while 290 indicated flat files, 147 for IMS, and 51 indicated Adabas.

Wayne informed the user group that data is growing at an outstanding rate. By some estimates, at least one and half to two zettabytes of data a year – to give you an idea, that is the equivalent of 250 billion DVDs or 36 million years of HD video.

Twitter’s own research in early 2012, estimated there’s roughly 175 million tweets every day.

Walmart handles more than 1 million customer transactions every hour, which are imported into databases estimated to contain more than 2.5 petabytes of data.

According to Forrester Research, the mobile payment market is predicted

to grow from \$12B to \$90B by 2017.

And given that mainframes are the gold standard for handling data, a good chunk of this data will end up being processed or stored on a mainframe.

As we know, mainframe databases and file structures are unique. What makes them so secure and so powerful also can make them challenging to access.

If you are a business analyst and your research includes mainframe data, you might encounter issues with compatibility. Most BI tools or analytics programs need data in a relational format, columns and rows, not a flat file or sequential file structure, and certainly not a hierarchical parent/child structure with multiple interdependencies.

In the past, the answer to



**Figure 3: Unique mainframe challenges**

mainframe data access was ETL. Extract, transform, and load, refers to the process for taking data from original sources – such as financial reports, shipments and revenue figures – and putting them into the right format to be loaded into a target such as a database or, more likely, a data mart or data warehouse. There trained professionals could submit queries and perform further analytics.

It was painstaking work. Standards were few and far between, data quality was often poor, and most ETL was done through custom-coding, perhaps with a team of C++ programmers. The process wasn't without problems – for example it

tended to adhere to siloed data practices and still didn't address data life-cycle issues – like how many copies of data were needed and how to keep a “single version of the truth”.

Accessing mainframe data through an off-host solution is a compromised solution:

- Off-host solutions are dedicated – they exist solely for the single purpose of either XML or SQL processing.
- Despite their single, dedicated purpose, off-host processing engines require their own separate hardware, network infrastructure, maintenance, management, power, and

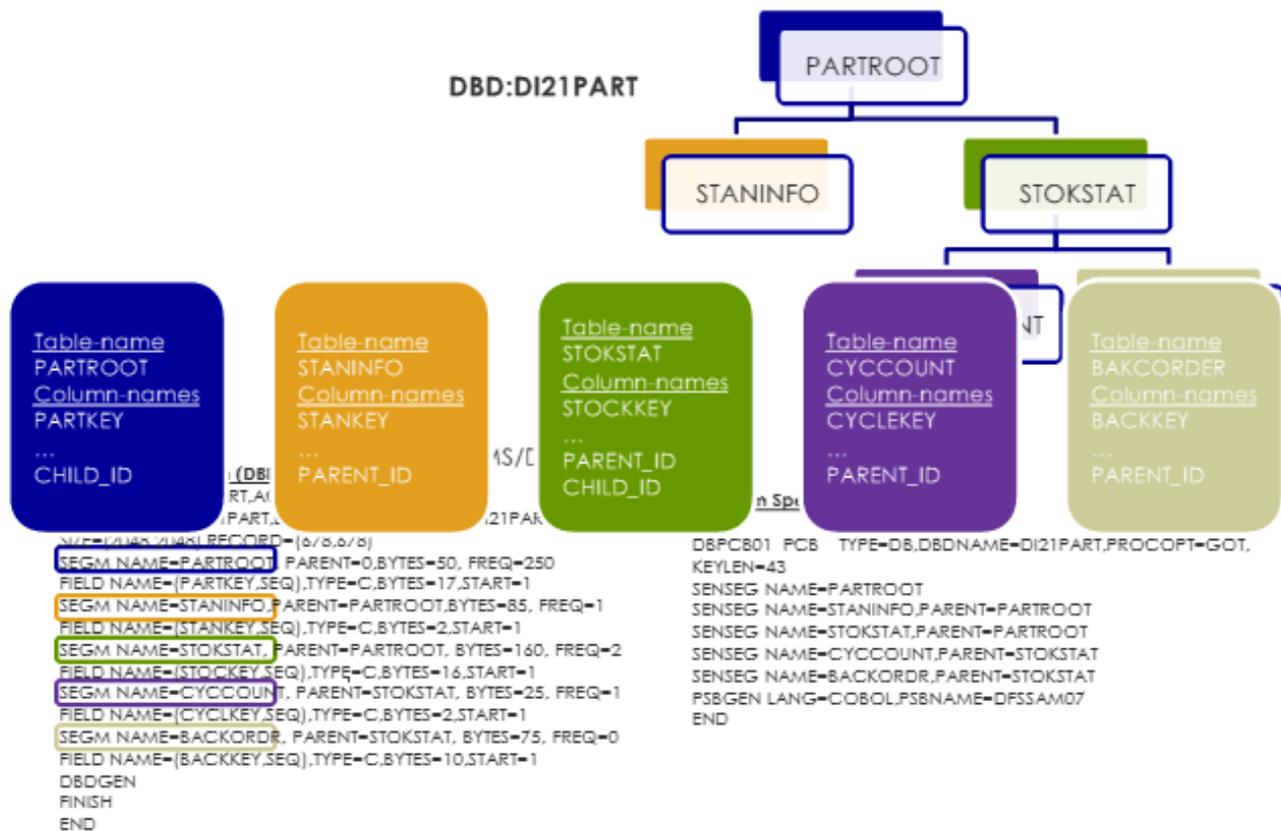
cooling – all adding to the cost.

- Off-host processing gateways don't share a coherent memory with the resources being accessed on the mainframe. Their use involves heavy I/O operations and the transfer of often enormous amounts of data across TCP/IP channels, resulting in inefficiencies and delays in operations, making real-time transactions impossible.

Prior to specialty engines, doing on-host data integration was too costly. The paradigm shift came with the zIIP.

zIIP can function as a fungible virtual appliance capable of processing both XML and SQL non-relational/relational transformation directly on the mainframe, leveraging the mainframe's processing power, but without incurring prohibitively steep mainframe system software consumption charges.

The zIIP being on the same coherent memory platform with the data and sharing the same Central Electronic Complex (CEC) is a tremendous advantage in terms of efficiency and performance. TCP/IP traffic is reduced to communication between



**Figure 4: DBD and PSB to virtual relational tables**

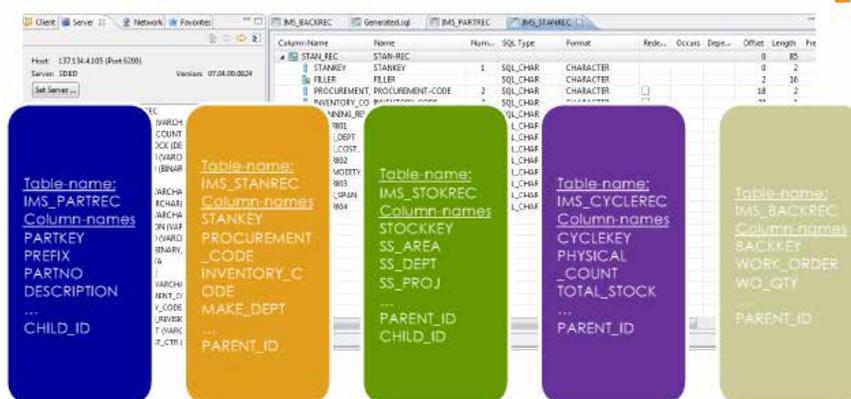
the requesting application and the mainframe source; the network hop between the mainframe and off-host processors is eliminated, as are all the affiliated I/O operations. The interaction between data source and processor appliance is a matter of memory-to-memory, with lightning fast transfer speeds in comparison to network transfers.

Wayne explained how DBDs and PSBs could be converted to virtual relational tables (see Figure 4). And how virtual relational tables could be created from

copybooks associated with DBD segments.

Not surprisingly, Wayne explained how Rocket Software's z/SQL software

could be used to achieve all this. He went on to show how SQL access to IMS/DB data was possible, and how all kinds of SQL joins could then be performed.



**Figure 5: Virtual relational tables from copybooks associated to DBD segments**

A copy of this presentation is available for download from the Virtual IMS user group Web site at [www.fundi.com/virtualims/presentations/zSQLAug13.pdf](http://www.fundi.com/virtualims/presentations/zSQLAug13.pdf).

You can see and hear the brief user group meeting by downloading the WMV file from [www.fundi.com/virtualims/presentations/2013-08-06meeting.wmv](http://www.fundi.com/virtualims/presentations/2013-08-06meeting.wmv).

### Meeting dates

- 8 October 2013, Glenn Witt, BMC Software will be discussing “Tuning Your Mainframe IMS Batch Processing”.

### IMS news

Parasoft has announced a new release of its Service Virtualization solution, which helps development and QA teams access the complete, realistic test environments needed to develop or test an application. They can capture real system behaviour by using monitors to record live transaction details, by analysing transaction logs, or by modelling behaviour from a simple interface. This latest release introduces advanced service virtualization for mainframes and enables fast, easy manipulation of complex data structures for

use in defining sophisticated simulation behaviour.

Parasoft’s Service Virtualization enables organizations to simulate the application under test’s interactions with subsystems including IMS, CICS regions, and DB2 databases. Communication may occur over protocols such as MQ, HTTP, TCP/IP, and JDBC, and it may involve payloads such as Copybook, XML, fixed-length or other custom message format. This not only promotes faster, more complete testing, but also minimizes disruption to high-value mainframe resources that could be better utilized for runtime or later-stage testing and validation. More information can be found at [www.parasoft.com/jsp/products/release.jsp?articleId=4361&type=Current](http://www.parasoft.com/jsp/products/release.jsp?articleId=4361&type=Current).

### Recent IMS articles

*Implementing IMS Synchronous Program-to-Program Switches Using DB2 for z/OS* by Edmond Torikian in *IBM Systems Magazine* (August 2013). You can find the article at [http://ibmsystemsmag.com/mainframe/administrator/ims/ims\\_db2\\_switch/](http://ibmsystemsmag.com/mainframe/administrator/ims/ims_db2_switch/).

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](http://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](mailto:trevor@itech-ed.com).

The Virtual IMS user group is free to its members.