



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, "MSU reduction due to in-memory table management with (any) IMS applications", and was presented by Gary Weinhold, a Systems Engineer with Data Kinetics, and Verna Bartlett, the Head of Marketing with Data Kinetics.

Gary has been working on IBM-compatible mainframe computers since the late 60s. His experience includes application programming in various environments, systems programming in MVS and VM, and software engineer and instructor for Amdahl Corporation. For the last twelve years, he has been a software developer for DataKinetics, working mainly on their tableBASE product. He has worked extensively

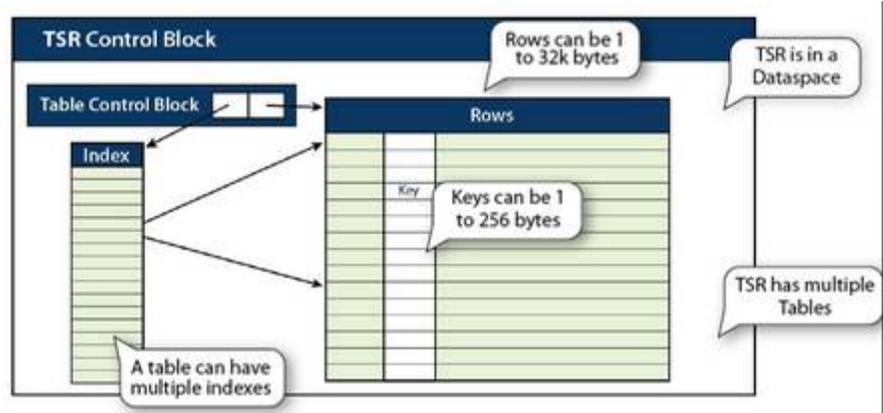


Figure 1: An in-memory table

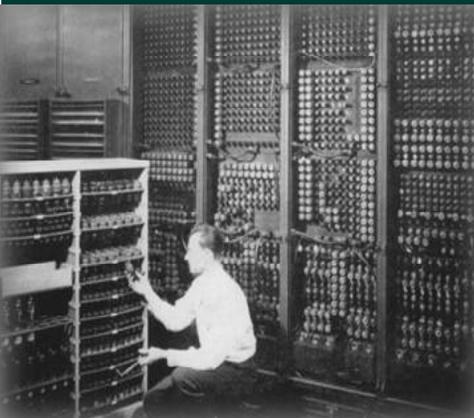
with customers, both on-site and remotely, to help them effectively use tableBASE in their environments. This has included training in tableBASE usage, suggestions for specific performance improvements in existing tableBASE applications, and recommendations for implementing tableBASE in non-tableBASE applications.

Verna has a broad range of marketing experience over a

wide range of products and industries. She has developed business plans, including target markets and positioning, for technology

Contents:

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

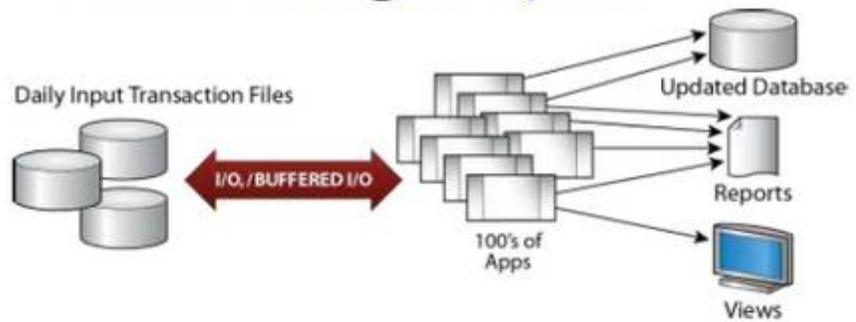


based start-ups in wireless, biosciences, health, and enterprise content management. She has developed the market strategy and overall calendar of marketing activities, which includes market launch plans and the supporting collateral development and marketing programs, for hardware and software products in public safety and telecom. She has developed channel marketing strategies, created and delivered channel education programs, including programs for sales teams and partners on a global basis (North America, Europe, Asia Pacific) for operational support systems in telecom. Prior to moving into marketing, she has been a technology manager for large data switches, both as a project manager and as a manager of software development teams and field support teams.

Their presentation started with an introduction to in-memory tables, moved on to what data works best in in-memory tables, described what you can do with in-memory tables, then examined the details of how in-memory tables can achieve its speed, and finally how in-memory tables solve critical business challenges.

Gary explained what was meant by an in-memory table, saying that with a Data Space in memory for table

Access data @ I/O speed



Access optimized @ memory speed

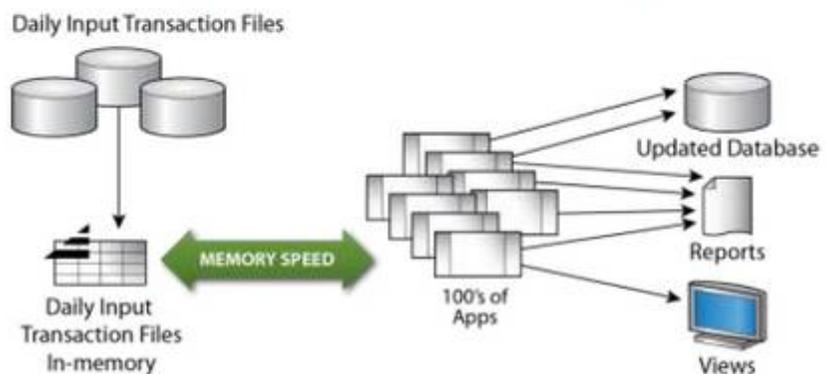


Figure 2: In-memory tables reduces elapsed time

processing, your in-memory tables are kept in a Table Share Region (TSR), and tables are made up of rows, keys, index, organization, and search methods (see Figure 1).

The speed advantages of using in-memory tables are shown in Figure 2. So, when not in memory, each request for data can result in one or more I/Os. IMS can buffer the data, but it still needs to be reformatted. IMS brings back a block, then the data is extracted and reformatted. With tableEXTENZ, the path to data is shorter and much,

much faster. The first call for data loads the table and creates an index. All other calls access in-memory tables, which return an entire row. It creates alternate indices on-the-fly, and virtual sorts use the alternate indices. It optimizes the search method, and there is no DBMS or OS overhead.

When the amount of time taken for each transaction is reduced, and the number of transactions per unit of time is increased, performance and capacity are increased significantly.

Gary informed us that our enterprises have many different kinds of data (see Figure 3):

- Some is used for transactions
- Some of it is temporary
- Some of it changes often, some changes infrequently
- We have all these different types of data, but we may have different naming conventions.

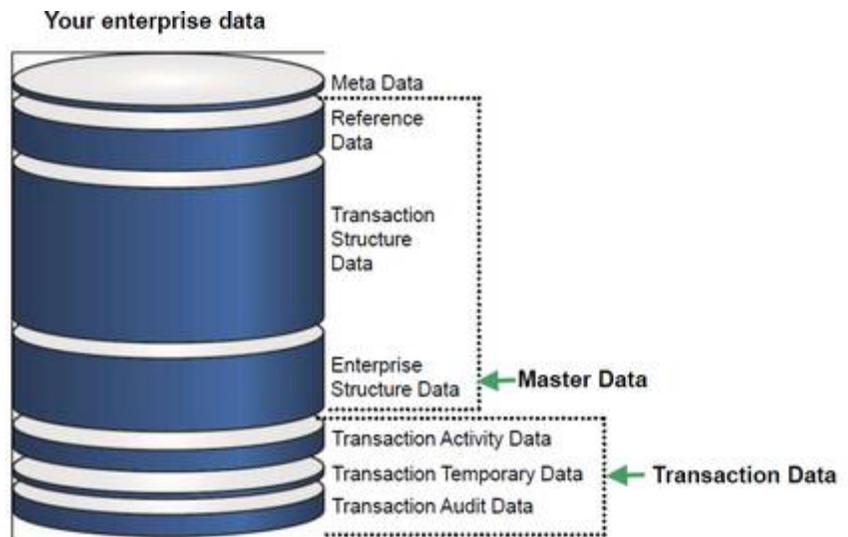


Figure 3: Enterprise data

Reference data:

- Is 5-15% of your total data
- Changes infrequently
- Is accessed often, and may represent as much as 80% of your accesses.

Temporary data:

- Is created, processed, and then deleted
- Generates a high volume of data accesses for the volume of data.

Remaining data:

- The largest volume of data
- Read often followed by a write
- The lowest number of accesses.

Reference data and temporary data benefit the most from optimization into in-memory tables. Reference data is data that is used to

validate transactions, or data that is looked-up as part of a transaction, or that categorizes other data.

In-memory temporary tables can be defined and used in lieu of DASD temporary tables. Temporary processing is done in memory and the table is deleted at the end of the processing. Tables can be shared between two or more programs or between two or more transactions.

In-memory tables improve performance because they do not do metadata translation on a field-by-field basis. They avoid OS access, eg link command, or I/O. They avoid GETMAINS and locking. they provide an efficient algorithm for search, choosing the search that is the most efficient for the data. They can return entire rows or portions

thereof. They use implicit commands to reduce changes to the calling application, and use shortcuts to reduce the pathlength to the data. And they allow the dynamic creation of indexes.

In-memory tables are also a good solution for batch processing because:

- Process large volumes, or batches, of instructions or files sequentially
- A large volume of work requires repetitive actions or repetitive logic
- Jobs often queued, with the output from one job providing the input for the next job
- Often have repeated reads of static data

With in-memory tables you can:	Decrease MSU	Decrease elapsed time	Increase flexibility and market adaptation	Reduce maintenance	Enable new paradigms
Place reference data in tables	Yes	Yes			
Replace temporary files with temporary tables	Yes	Yes			Yes
Use tables for rules			Yes	Yes	
Use tables as a message queue	Yes	Yes			Yes
Use decision tables to replace logic trees	Often		Yes	Yes	
Use tables for process control			Yes	Yes	
Use temporary tables for implementing complex algorithms	Yes	Yes	Yes	Yes	Yes

Figure 4: In-memory table results

- May create temporary files as part of the transaction processing
- Batch jobs do a lot with reading, writing, and sorting sequential and VSAM files
- Manual intervention or input not required
- Runs automatically once a batch job begins, it continues until it is done or until an error occurs
- Traditionally there has been a batch window – a period of time when there would not be any OLTP, and batch jobs would run uninterrupted.

Verna summarized the results that in-memory tables can provide – see Figure 4..

A copy of Gary and Verna’s presentation is available for download from the Virtual IMS

user group Web site at www.fundi.com/virtualims/presentations/in-memory_tablesJun11.pdf.

You can see and hear the whole user group meeting by downloading the WMV file from www.fundi.com/virtualims/presentations/2011-06-14meeting.wmv.

Meeting dates

The following meeting dates have been arranged for the Virtual IMS user group:

- 9 August 2011 – Scott Quillicy, SQData. “IMS replication for high-availability”.
- 11 October 2011 – Rosemary Galvan, Principal Software Consultant – IMS, BMC. “Database Performance – Could Have, Should Have, Would Have”.

- 6 December – IBM presenter, “IMS 12”.

Recent IMS articles

Best Practices: IMS High Availability Large Database Maintenance by Sudipta Sengupta in *z/Journal* (June/ July 2011). You can find the article at www.mainframezone.com/it-management/best-practices-ims-high-availability-large-database-maintenance.

IMS news

NEON Enterprise Software, which recently lost its court battle with IBM over zPrime, has sold its IMS portfolio to BMC Software. Full information can be found at www.bmc.com/news/press-releases/2011/bmc-acquires-ims-assets-of-neon-enterprise-software.html.

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.