

Variable Size Data Pages – Understanding the Pros and Cons

October 15 – 17, 2001



Abstract

Ingres II supports multiple sizes of data pages. This presentation provides a technical overview of this feature, provides good examples for use of this feature, and explores some of the issues of implementation of this feature. This is not intended to be an all-inclusive list of features, nor is it intended to be a step-by-step implementation guide. Rather, it is intended to provide some important basic information to assist with the decision-making process on whether or not this new feature will benefit your site. This can be a very powerful feature when used properly and hopefully this presentation will help provide a better understanding of the benefits of variable size data pages.



Biography

- Chip Nickolett
- Comprehensive Consulting Solutions, Inc.

Chip is the founder of Comprehensive Solutions (www.Comp-Soln.com), was a Consultant and Consulting Manager for Computer Associates, and was a Senior Consultant for the Ingres Products division of the ASK Group. Chip is also a past President of the North American Ingres Users Association (www.naiua.org).



Topics of Discussion

- Overview
- What Page Size is Best?
- How Many Different Page Sizes Do I Need?
- Benefits of Larger Pages
- Limitations of Larger Pages
- Other Considerations
- Summary



What Page Size is Best?

- There are no single correct answer for this question!



What Page Size is Best?

- There are no single correct answer for this question!
- Are larger row (tuple) sizes required?



What Page Size is Best?

- There are no single correct answer for this question!
- Are larger row (tuple) sizes required?
- **Is data page utilization poor due to row size or page size?**



What Page Size is Best?

- There are no single correct answer for this question!
- Are larger row (tuple) sizes required?
- Is page utilization poor due to row size?
- **Is row-level locking required?**



What Page Size is Best?

Sample spreadsheet showing the baseline number of pages for a table for each Page Size

Table Name	Rowsize	Est # of Rows	#2K pages	#4K pages	#8K pages	#16K pages	#32K pages	#64K pages
ord_stat	5	100000	250	776	384	196	196	196
ord_hdr	1005	100000	100000	33334	14286	6667	3226	1588
ord_dtl	158	650000	54167	30953	14773	7387	3673	1831
ord_text	500	50000	12500	7143	3334	1667	807	404



How Many Different Page Sizes Do I Need?

- Again, there is no single correct answer!



How Many Different Page Sizes Do I Need?

- Again, there is no single correct answer!
- **Generally 1 or 2 different page sizes are sufficient**



How Many Different Page Sizes Do I Need?

- Again, there is no single correct answer!
- Generally 1 or 2 different page sizes are sufficient
- **Remember, tables and their associated indexes are not required to share the same page size (2 KB index pages are generally best)**



How Many Different Page Sizes Do I Need?

- Again, there is no single correct answer!
- Generally 1 or 2 different page sizes are sufficient
- Remember, tables and their associated indexes are not required to share the same page size
- **Use more when it makes sense!**



Benefits of Larger Data Pages

- **Provide the ability to have much larger tuple sizes**



Benefits of Larger Data Pages

- Provide the ability to have much larger tuple sizes
- **Have the potential to maximize data utilization within a table**



Benefits of Larger Data Pages

- Provide the ability to have much larger tuple sizes
- Have the potential to maximize data utilization within a table
- **Have the potential to minimize I/O (thereby improving performance)**



Benefits of Larger Data Pages

- Provide the ability to have much larger tuple sizes
- Have the potential to maximize data utilization within a table
- Have the potential to minimize I/O (thereby improving performance)
- **Can make the table smaller in size when compared to 2 KB pages**



Benefits of Larger Data Pages

- **Provides the ability to use row-level locking**



Benefits of Larger Data Pages

- Provides the ability to use row-level locking
- **Provides the ability to use the “Alter Table” statement for schema modification**



Benefits of Larger Data Pages

- Provides the ability to use row-level locking
- Provides the ability to use the “Alter Table” statement for schema modification
- **DMF cache can be tuned differently for each page size**



Benefits of Larger Data Pages

- Global temporary tables can utilize different DMF cache pools



Benefits of Larger Data Pages

- Maximum tuple size increases for the larger page sizes
 - 2 KB = 2008 bytes
 - 4 KB = 3988 bytes
 - 8 KB = 8084 bytes
 - 16 KB = 16276 bytes
 - 32 KB = 32660 bytes
 - 64 KB = 32767 bytes



Limitations of Larger Data Pages

- Data may be row limited by the TID (tuple identifier)



Limitations of Larger Data Pages

- Data may be row limited by the TID (tuple identifier)
- **Maximum number of rows per page**
 - » 2 KB = 511
 - » 4 KB = 138
 - » 8 KB = 279
 - » 16 KB = 512
 - » 32 KB = 512
 - » 64 KB = 512



Limitations of Larger Data Pages

- **Increased overhead for both the page and row can limit utilization**
 - 80 bytes for the page
 - 26 bytes for each row



Limitations of Larger Data Pages

- Increased overhead for both the page and row can limit utilization
- **Potential to have a negative impact on performance if improperly configured due to increased I/O rates**



Limitations of Larger Data Pages

- Increased overhead for both the page and row can limit utilization
- Potential to have a negative impact on performance if improperly configured due to increased I/O rates
- **Overall table size may increase when compared to 2 KB pages**



Other Considerations

- **Optimizer behavior and QEPs in your environment may be different**
 - Optimizer algorithms are supposed to be the same regardless of page size
 - Number of pages, especially with group buffers, can skew the plan



Other Considerations

- Optimizer behavior and QEPs in your environment may be different
- **I/O performance on your hardware**
 - Disk block or stripe size
 - RAID



Other Considerations

- Optimizer behavior and QEPs in your environment may be different
- I/O performance on your hardware
- **Locking system reconfiguration is usually necessary when row-level locking is used**
 - Set transaction isolation level to read committed for BTREE tables



Other Considerations

- Will the various DMF caches share the same memory segment?
 - The size of shared memory segments is 2 GB on most systems



Other Considerations

- Will the various DMF caches share the same memory segment?
- Available memory (RAM) on your hardware
 - Consider that a 15,000 page cache for 64 KB data pages is almost 1 GB in size



Session Summary

- **Variable Size Data Pages can be a very useful feature...**
 - when larger tuple sizes are required
 - to minimize I/O (and increase performance)
 - to decrease the size of a table through better utilization of the data page
 - when selectively applied to tables and indexes
 - for row-level locking



Session Summary

- **It is important to remember the limitations...**
 - page is potentially row limited by the TID
 - has the potential to increase the table size



Session Summary

- It is important to remember the limitations...
- **Test various configurations in your own specific environment**



Session Summary

- It is important to remember the limitations...
- Test various configurations in your own specific environment
- **Gather performance metrics, review QEPs, compare results and make an informed choice!**



Session Summary

The more you understand this feature the better you will be able to use it to your benefit



Questions & Answers

