

Ten Clicks to Breakthrough DB2 UDB Results

Philip K. Gunning
Performance Consultant
DGI
Session: D9

October 23, 2002
11:00 AM – 12 Noon

© *Copyright 1999*



**INTERNATIONAL
DB2 USERS GROUP**



Independent • Not-for-Profit • User Run

ABSTRACT

- This session will present key DB2 UDB Performance Measurements and a path to navigate through the numbers such that huge performance gains can be achieved very rapidly, often in just minutes. Discover the top warning signs and symptoms of poor database health, and learn how to drill down to find the real underlying problems.



OUTLINE

1. 10 Key Performance Tips
2. Key Performance Knobs
3. Memory Utilization -- Agent parameters
4. Evaluating Sort Performance
5. Monitoring and Measuring results



DISCLAIMERS

DB2™ is a registered trademark of IBM Corporation. Oracle is a registered trademark of Oracle Corporation and SQL Server is a registered trademark of Microsoft.

DB2 Universal Database (DB2 UDB) is a registered trademark of International Business Machines Corp.



Other logos and product/trade names are registered trademarks or trademarks of their respective companies.

- **BREAKTHROUGH DEFINITION:**

NOUN. “ MAJOR ADVANCE OR DISCOVERY”



NOUN.

“ ACT OF BREAKING THROUGH AN OBSTACLE”

The Ten Clicks

1. [SQL](#)
2. [Indexes](#)
3. [Bufferpools](#)
4. [Prefetching](#)
5. [Parallelism](#)
6. [SORT](#)
7. [Logging](#)
8. [AGENTS](#)
9. [Memory Structures](#)
10. [Monitoring](#)



[Continue](#)

SQL



SQL

- Poor performing SQL is usually 70 - 80% of most performance problems
 - Using the 80- 20 rule, 20% of this SQL will be what you need to focus your efforts on
- 
- Just because it is a Third Party Vendor Package doesn't mean YOU CAN'T TUNE IT!
 - Index changes
 - Clustering indexes*
 - Include columns

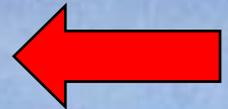
SQL

- db2 get snapshot for dynamic sql on hrproddb
- select STMT_TEXT from
table(SNAPSHOT_STATEMENT)'ATGW_D', -2)) as D;
 - Returns statement information from application snapshot
- db2 create event monitor dgi_han for statements write
to table buffersize 4 blocked manualstart ;
- db2advis -d atgw_p -i inputfile.txt



SQL

- ✓ TIP: Identify SQL executed with high frequency and tune. Review, monitor, and tune dynamic and static SQL on a regular basis



INDEXES



Indexes

- New Type-2 Index in V8.1
 - Eliminates most next key share locking problems
 - Pseudo delete of index entries
 - **REQUIRED** for MDC
 - Migration Considerations
 - Migrate tables at-a-time
 - Type-2 required for New V8.1 Online Reorg and Online Load
 - Cannot mix and match type-1 and type-2 on same table
- 

Indexes

- Clustering
 - Maintains data rows in order of index entries
 - Greatly improves prefetching and sequential scans
 - Every table should have a clustering index
 - Clustering reduces or eliminates sorts
 - Most Leading ERP and CRM Third Party Vendor packages do not use any clustering indexes
 - Must reorg table using clustering index for clustering to take effect
 - One clustering index per table
- 
- A fighter jet is flying in the sky, positioned behind the text of the second bullet point. The jet is a dark color, possibly black or dark grey, and is shown from a side profile, flying towards the right. The background is a clear blue sky.

Indexes

- Read-only objects can have more indexes than heavily inserted and updated tables
 - Review application access patterns with developers before indiscriminately adding indexes



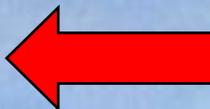
- Consider “include columns” for index-only access
- Reorg indexes
- Index Anding and Index Oring

Indexes

- Multidimensional Clustering (MDC)
 - can create indexes on multiple dimensions
 - no reorg required, automatically kept in clustering order
 - stored in blocks
 - MDC enables tables to be clustered on multiple keys or dimensions, simultaneously
 - primarily designed to benefit DW/BI type queries but can also be used for OLTP
- 
- A fighter jet is shown in flight against a blue sky background, positioned behind the text of the list.

Indexes

- ✓ Create appropriate indexes
- ✓ Consider include columns for index-only access
- ✓ Use clustering indexes to improve performance
- ✓ Analyze using Index Advisor and look at local, order by and join predicates



BUFFERPOOLS



Bufferpools

- Asilomar Report
 - 1998 meeting of Relational Database Industry and Academic Professionals
- Predict main memories and addressing techniques will improve to hold vast amounts of data
- Indicate Significant improvement needed in Optimization Technology
- Sizing
 - 75% of RAM available if dedicated database server
 - Better to over-allocate than to under-allocate
 - monitor OS paging and reduce bufferpool size if necessary



Bufferpools

- Breakout Strategy
 - Random versus sequential
 - Index versus data
 - Prioritization
 - Large Randomly Accessed Tables
 - DW versus OLTP
 - One large bufferpool
 - Reduces overhead of managing multiple pools
 - DW tables typically very large involving large scans
- 

Bufferpools

- ✓ Allocate up to 75% of memory to bufferpools
- ✓ Use block bufferpools for large sequential access type applications
- ✓ Use one bufferpool for DW/BI
- ✓ Use multiple bufferpools and associated breakout strategy for OLTP



PREFETCHING



Prefetching

- Num_ioservers should be set to one-half the number of containers in the database as a starting point
 - monitor and adjust as necessary
- Exentsize of 32 is generally good
 - If set too high can overwhelm bufferpool
- Prefetchsize should be set to (extentsize x number of containers)
 - This will enable parallel prefetch request to be issued to each container



Prefetching

- Parallel IO
 - DB2_PARALLEL_IO=* (1,2,3)
 - DB2_STRIPED_CONTAINERS
 - RAW DEVICES ARE BECOMING THE NORM!
 - Prefetchsize= (extent size x # of containers)
- Spread containers over multiple physical disks
- Set Operating System parameters to maximize Asynchronous IO performance (AIX - minservers and maxservers)
- Block-based bufferpools in V8.1 provide improved performance by prefetching sequential pages that are contiguous in block bufferpools



Prefetching

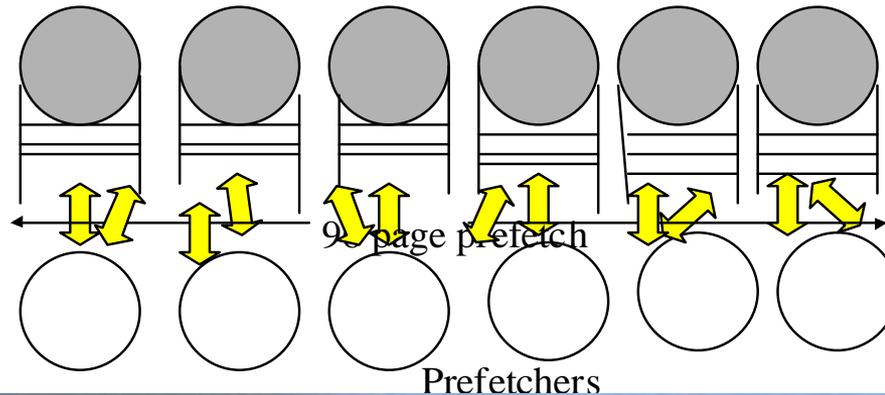
RAID-5 (6 + 1) EXAMPLE

Physical Disk Striping

- Each extent spread across two disks
- Single prefetch drives all 6 drives simultaneously

but requires two IOs from each, instead of 1!

```
Db2set DB2_STRIPED_CONTAINERS=ON  
Db2 create tablespace New_GUY  
Managed by database  
Using (Device 'prdrhd1' num_pages)  
Pagesize 4k  
Extentsize 16 prefetchsize 96  
Db2set DB2_PARALLEL_IO=*
```



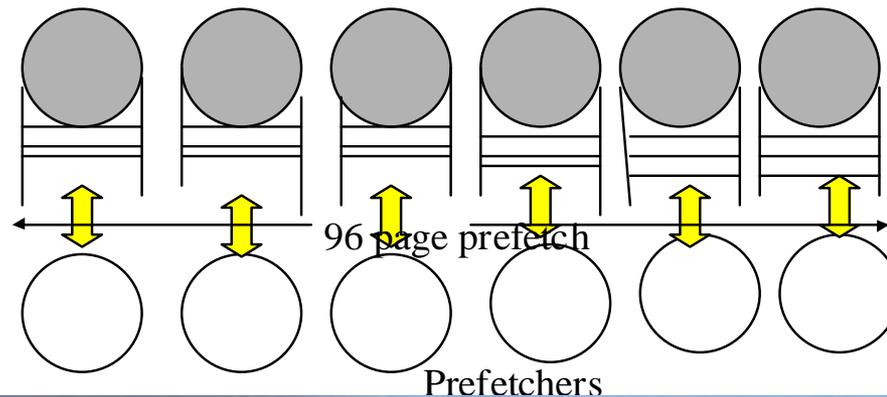
Prefetching

RAID-5 (6 + 1) EXAMPLE

Physical Disk Striping

- Each extent aligned on an internal stripe boundary on RAID-5 Device
- Single prefetch drives all 6 drives simultaneously
- Each disk must handle only 1 IO

```
Db2set DB2_STRIPED_CONTAINERS=ON
Db2 create tablespace New_GUY
Managed by database
Using (Device 'prdrhd1' num_pages)
Pagesize 4k
Extentsize 16 prefetchsize 96
Db2set DB2_PARALLEL_IO=*
```



Prefetching

- ✓ Set prefetchsize to a multiple of extentsize x number of containers
- ✓ Don't overestimate prefetchsize as this can overwhelm bufferpool
- ✓ Monitor Asynchronous Pages Read per Minute to evaluate efficiency of prefetching



PARALLELISM



Parallelism

- OLTP versus DW
 - OLTP set `intra_parallel` to NO
 - DW set `intra_parallel` to YES
 - How to disable for OLTP
 - set `MAX_QUERYDEGREE=1`
 - set `DFT_DEGREE=1`
 - How to Exploit for BI/DW
 - set `DFT_DEGREE=1` and let DB2 decide
 - set `MAX_QUERYDEGREE` equal to number of CPUs but never more
- 
- A fighter jet is flying in the sky, positioned behind the text of the list. The jet is a dark color, possibly black or dark grey, and is shown from a side profile, flying towards the right. The background is a clear blue sky.

Parallelism

- Utility Parallelism
 - Utilities use parallelism independently of DB or DBM CFG settings
 - New “throttle” capability in V8.1
 - Fast Communication Manager (FCM)
 - used to manage communication of messages between subagents
 - used with `intra_parallel` enabled or if using ESE Partitioning Option
 - Monitor using “Get snapshot for fcm on all nodes”
- 

Parallelism

- ✓ Set `intra_parallel` to `no` in OLTP
- ✓ Use utility parallelism with `load`, `reorg`, `backup` and `restore` (Ensure adequate buffer size)



SORT



Sort

- Private
 - Sortheap allocated out of agent private memory
- Shared
 - Sortheap allocated out of Database Global Memory
- SHEAPTHRES
 - Private Sort – Soft Limit
 - Shared Sort – Hard Limit



Sort

- Piped
- Non-Piped
- Sort Overflows
 - Entire sort overflows
 - Overflows to Bufferpool and System Temporary Space Work Tables
- How to Avoid



Sort

- Post Threshold Sorts
- Piped Sorts Rejected
- Hash Joins
- Binary Sort



- DB2_BIN_SORT in V6.1 enables new sort algorithm that optimizes sort

Sort

- Eliminate through
 - Clustering
 - Indexing
 - Efficient SQL



- Cannot exterminate because we humans need things returned in an orderly fashion, but we can eliminate the unnecessary ones!

Sort

- ✓ SORT IS A FOUR LETTER WORD
- ✓ Find the SQL that is performing sorts and eliminate through indexing, clustering, changing the SQL
 - ✓ Use Snapshots, Event Monitors, Third Party Vendor tools, DB2 Explain



- ✓ Index Advisor
- ✓ MDC



LOGGING



Logging

- Primary Logs formatted at first connect or activate database command
 - Have enough primary and plenty of secondary
 - Determine via benchmarking and monitoring
- Maximize use of log buffer by setting logbufsz to at least 64 4k pages (32-bit AIX)
- Increases throughput and makes maximum use of log buffering



Logging

- Mincommit
 - Used to group commits to increase throughput
 - Can increase throughput at expense of response time wait up to 1 second
- Create separate file system for logs and don't put any other data on volumes with logs
- Consider mirroring logs or use newlogpath2 DB CFG parameter



Logging

- ✓ In an OLTP environment set mincommit using following formula
(# of transactions per second / 10) and use the result for setting mincommit.



- ✓ Example: 95 TX per sec / 10 = mincommit 9
- ✓ Set logbufsz to at least a value of 64, larger if ample memory available



AGENTS



Agents

- Coordinator Agent – Each application has one which does work on its behalf and in a parallel environment distributes work to subagents
- Upon disconnect or detach from an instance the coordinating agent is freed and marked as idle if max number of pool agents not reached else it is terminated and storage freed if max number of pool agents reached
- DBM CFG parameter max_coordagents



Agents

- Maximum Number of Agents (maxagents)
 - specifies the maximum number of database manager agents, whether coordinating agents or subagents, available at any given time to accept application requests



- Can be used in resource constrained systems to limit memory usage

Agents

- Maximum Number of Active Applications (maxappls)
 - Specifies the maximum number of concurrent applications that can be connected to a database
 - When reached, an error is returned to the application and connection is not allowed
- Can be used to throttle applications in a resource constrained system



Agents

- Maximum Number of Concurrent Agents (maxcagents)
 - Specifies the max number of database manager coordinating agents that can be concurrently executing a database manager transaction



- Does not limit the number of applications connected but limits the number of database manager agents that can be processed concurrently
- Can be used to throttle applications if resource constrained

Agents

- Initial Number of Agents in Pool (Num_initagents)
 - Specifies the initial number of idle agents that are created in the agent pool at DB2START



- By specifying a value, agents are available in the pool for initial requests and the overhead of repeated agent creation is avoided

Agents

- Agent Pool Size (num_poolagents)
 - Specifies how large the agent pool can get
 - Contains subagents and idle agents
 - Idle agents can be used as coordinating agents or subagents
 - If more agents created than this parameter they will be terminated when the current request is completed rather than returned to the pool



Agents

- ✓ Ensure that there are no Stolen Agents!
- ✓ Ensure that no Agents are waiting for a Token (maxcagents)
- ✓ Use connection pooling in OLTP and hybrid OLTP/DW environments
- ✓ Connection concentration
 - ✓ Reduces amount of total memory used by agents and maps many connections to one logical agent



MEMORY



Memory

- Locklist
 - Used to store memory for locks
 - Over-allocate versus under-allocate
 - num_locks default of 22% is a good starting point
 - Use lock_timeout to set timeout for lock waiters
 - Default -1 means applications wait for lock forever
 - not normally a good thing
 - Catalog Cache
 - Used to store table descriptors
 - Eliminates read to catalog if used efficiently
 - Size and monitor to ensure consistent hit ratio > 90%
- 
- A fighter jet is flying across the middle of the slide, partially overlapping the text. It is a dark-colored aircraft with a delta-wing configuration, flying from left to right against a clear blue sky.

Memory

- Package Cache
 - How it really works
 - Must be large enough in the first place to store all executing dynamic and static SQL before caching occurs
 - Used to store dynamic and static SQL
 - Size and monitor appropriately and make sure inserts are low
 - If package cache overflows, it overflows into locklist and other dbheap areas!
 - Size and monitor to ensure hit ratio > 90%
 - New V8.1 capability to flush package cache!
- 
- A fighter jet, possibly an F-16, is shown in flight against a clear blue sky. The jet is positioned horizontally across the middle of the slide, with its nose pointing to the right. The image is slightly blurred, suggesting motion. The jet is white with some dark markings on the nose and wings.

Memory

- ✓ Over-allocate locklist by 50% or ensure that no more than 50% of locklist is in use
- ✓ Make catalogcache_sz large enough so that a 90% hit ratio is obtained
- ✓ Sufficiently size bufferpools!
- ✓ Analyze memory usage over time



MONITORING



Monitoring

- New SQL Function in V8.1 enables taking snapshots via SQL
- `select STMT_TEXT from table(SNAPSHOT_STATEMENT)'ATGW_D', -5)) as T;`
- Can issue from command line or API
- Event monitors can write to tables!



Monitoring

- Snapshot monitoring enabled via DBM CFG monitor switches or application can set switches
- Snapshot capability added to SQL as function
- Snapshot output on LOCKS improved to show lock details and SQL involved!



Monitoring

- Event Monitoring can now be created in EEE and write to a table on each node!



Monitoring

- New V8.1 Health Monitor enables alerts by exception
 - send email to predefined list of DBAs
 - send alert to a pager
- Exception monitoring fairly common requirement with many organizations having 100's or thousands of databases to monitor
 - many companies had written in-house exception monitoring scripts but were still looking for a better solution
- Health monitor is configured via the Health Center



HEALTH CENTER

The screenshot shows the Health Center application window. The title bar reads "Health Center". The menu bar includes "Health Center", "Selected", "Edit", "View", "Tools", and "Help". The toolbar contains various icons for navigation and actions, along with a refresh button and a "1 minute" refresh interval dropdown. The left pane shows a tree view with "DB2" and "TOOLSDB" under it. The main pane displays a table of alerts for "DB2 - TOOLSDB".

Health Indicator	Value	Category	Partition	Object Type	Timestamp of Issuing Alert
DB2 - TOOLSDB					
Alarm					
Catalog Cache Hit Ratio	65	Package and C...	0	Database	2002-07-24 20:32:42.148
Warning					
Package Cache Hit Ratio	72	Package and C...	0	Database	2002-07-24 20:32:42.148

2 of 2 items displayed

By Object Name, Severity View

Monitor data is refreshed. (7/24/02 8:39 PM)

HEALTH MONITOR

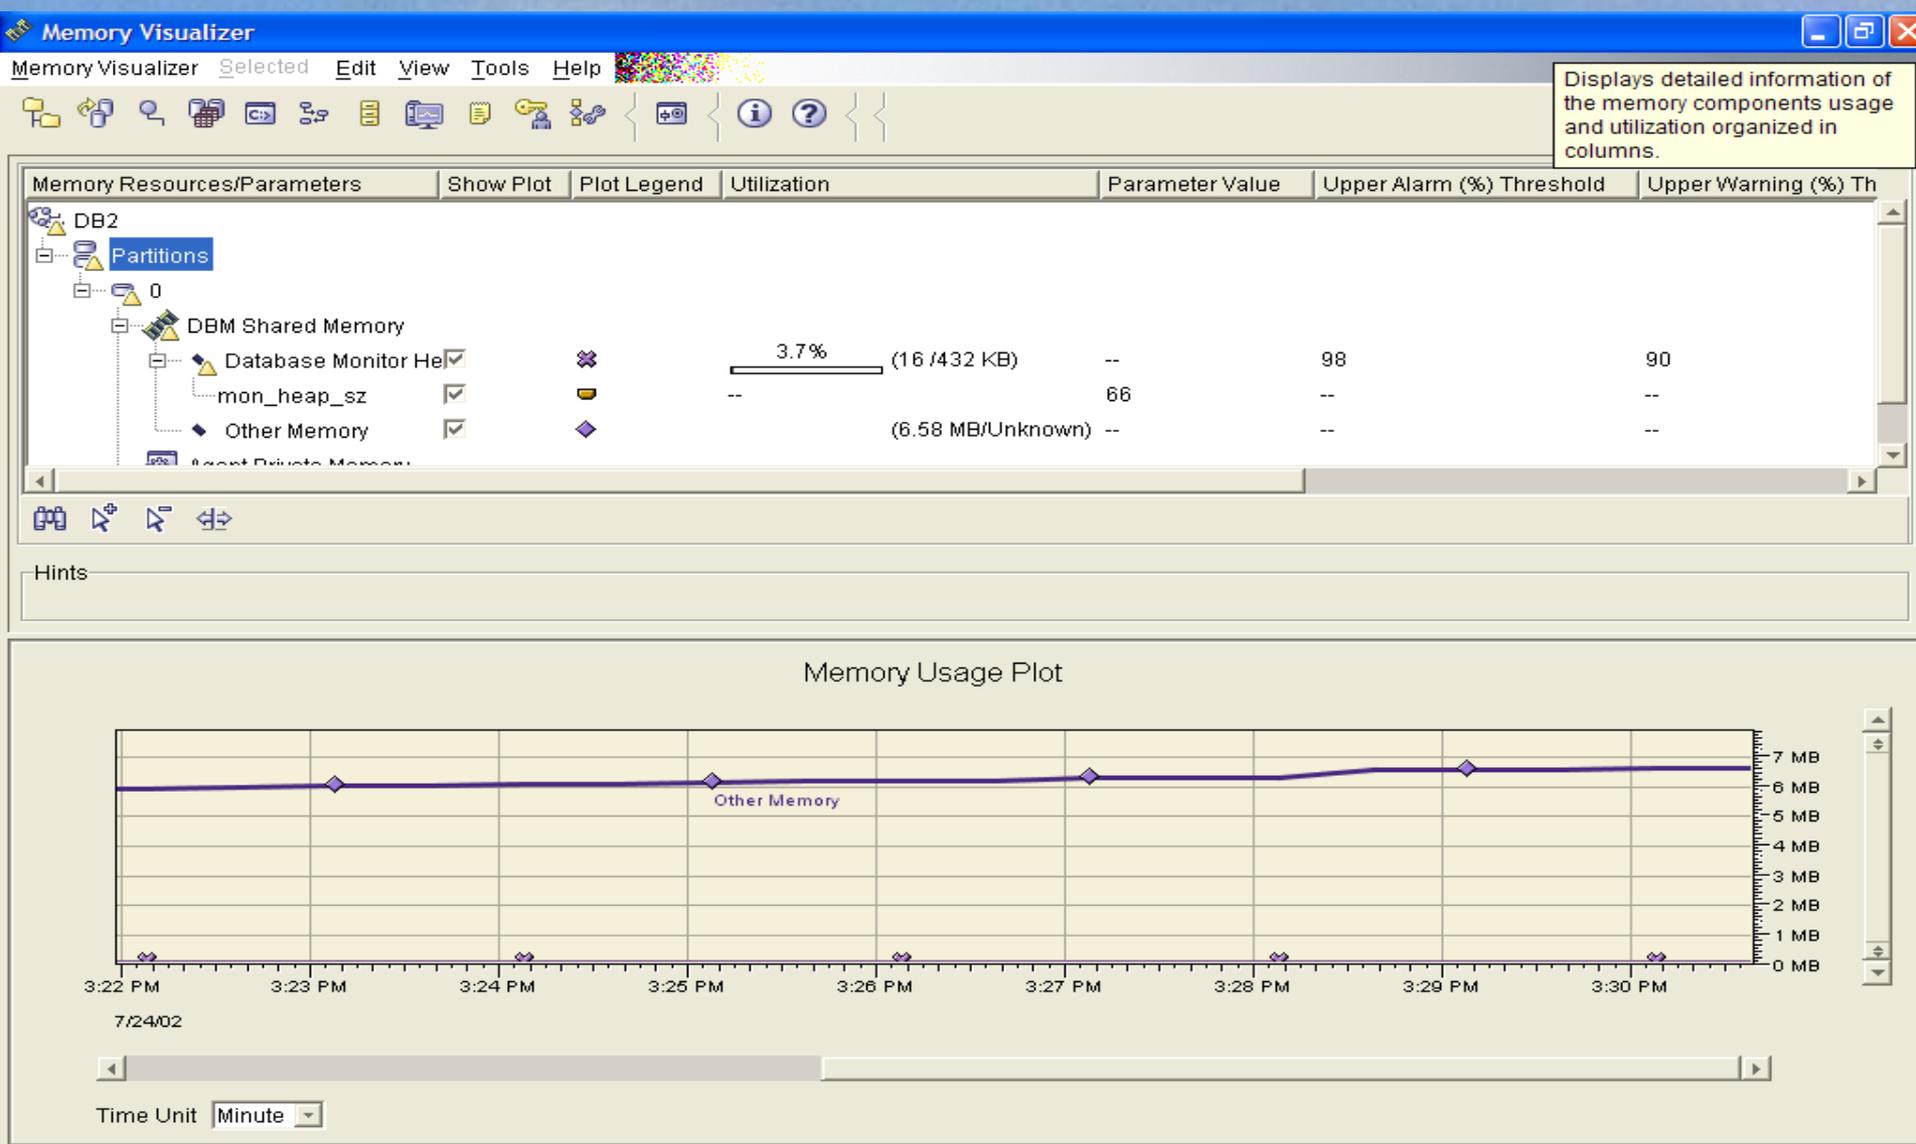
- Server side code that monitors thresholds and generated alerts via Health Center, Beacons, Email, page to pager
- Configured via the Health Center



- Implements “Management by Exception” alerting
 - It address the need to monitor 100s or 1,000s of databases without “eyes on”



MEMORY VISUALIZER



Monitoring

- ✓ Have a monitoring plan and monitor on a regular basis
- ✓ Take advantage of new monitoring tools in V8.1 and from Third Party Vendors
- ✓ If you are not monitoring, you are driving blind



Summary

- Now you have 10 ways to improve your database performance
- Remember, make one change at a time and monitor and measure the effects of all changes



References

- SC09-4848-00, What's New, V8
 - SG24-6012-00, Redbook, DB2 UDB V7.1 Performance Tuning Guide
 - SC09-2945-01, DB2 UDB V8.1 Administration Guide: Performance
- 
- A fighter jet, possibly an F-16, is shown in flight against a blue sky. The jet is viewed from a low angle, showing its wings and tail. It appears to be moving from left to right across the frame.
- Asilomer Report, ACM Proceedings, 1998
 - SC09-2946-01, DB2 UDB V7.2 Administration Guide: Planning
 - AIX Performance and Tuning Guide

10 Clicks to Breakthrough DB2 UDB Results
Session: D9

Philip K. Gunning
DGI

THANK YOU!

**pgunning@breakthroughdb2.
com**

© Copyright 1999

**INTERNATIONAL
DB2 USERS GROUP**



Independent • Not-for-Profit • User Run