

Tuning Your DB2 Buffer Pools

How Much Memory do You Really Need?

Responsive Systems Company

281 Highway 79

Morganville, NJ 07751

www.responsivesystems.com

(732) 972-1261

How much memory do you need?

- **The current “bar” on 64-bit machines**
 - Maximum available today is 128 Gigabytes
 - It will go up...
- **But, how much do you need for your system?**
 - We always want more
 - We have always wanted more
 - We have always filled it up
- **There are a “few” huge systems that may benefit from massive pool sizes**
 - Nobody is going to push this current bar for a while

How much memory do you need?

- How much are you using now?
 - How much more do you think you need?
- Are you sure that more memory will provide better performance?
- We will see today that “more” does not always provide better performance

Who is on a 64-bit Machine now?

- **DB2 V8 installed?**
- **DB2 V8 ordered?**
- **DB2 V8 Planned?**
 - This year?
 - Next year?
 - Someday?

Will more memory help *You?*

- How can you find out?
- Play the guessing game...
 - Just throw it at the system and pools
- But will you get a real benefit?
- Measure before/after performance?
 - When will you measure it?
 - How will you measure it?
 - What is really important as a performance metric?

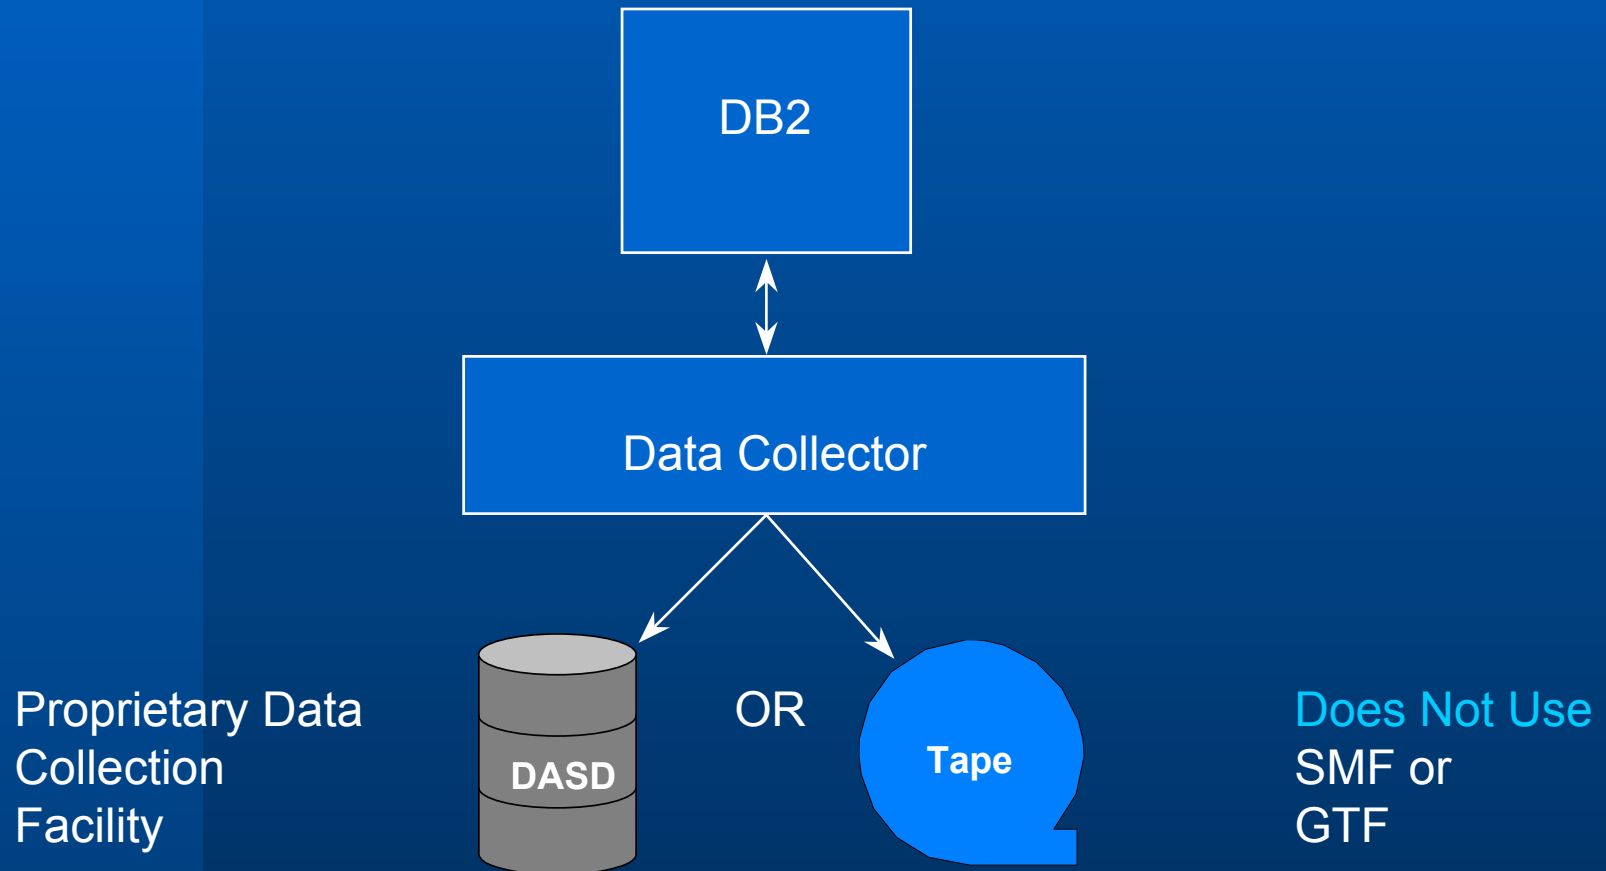
Buffer Pool Tool® for DB2

- Is the only product that can **Predict** the effect of changes to the most important tuning metric – *the I/O rate/second*
- Provides detailed statistical performance data
- Provides a **proven methodology** for pool tuning
- Shows you how to get it right the first time
- Leverage your system resources while improving performance

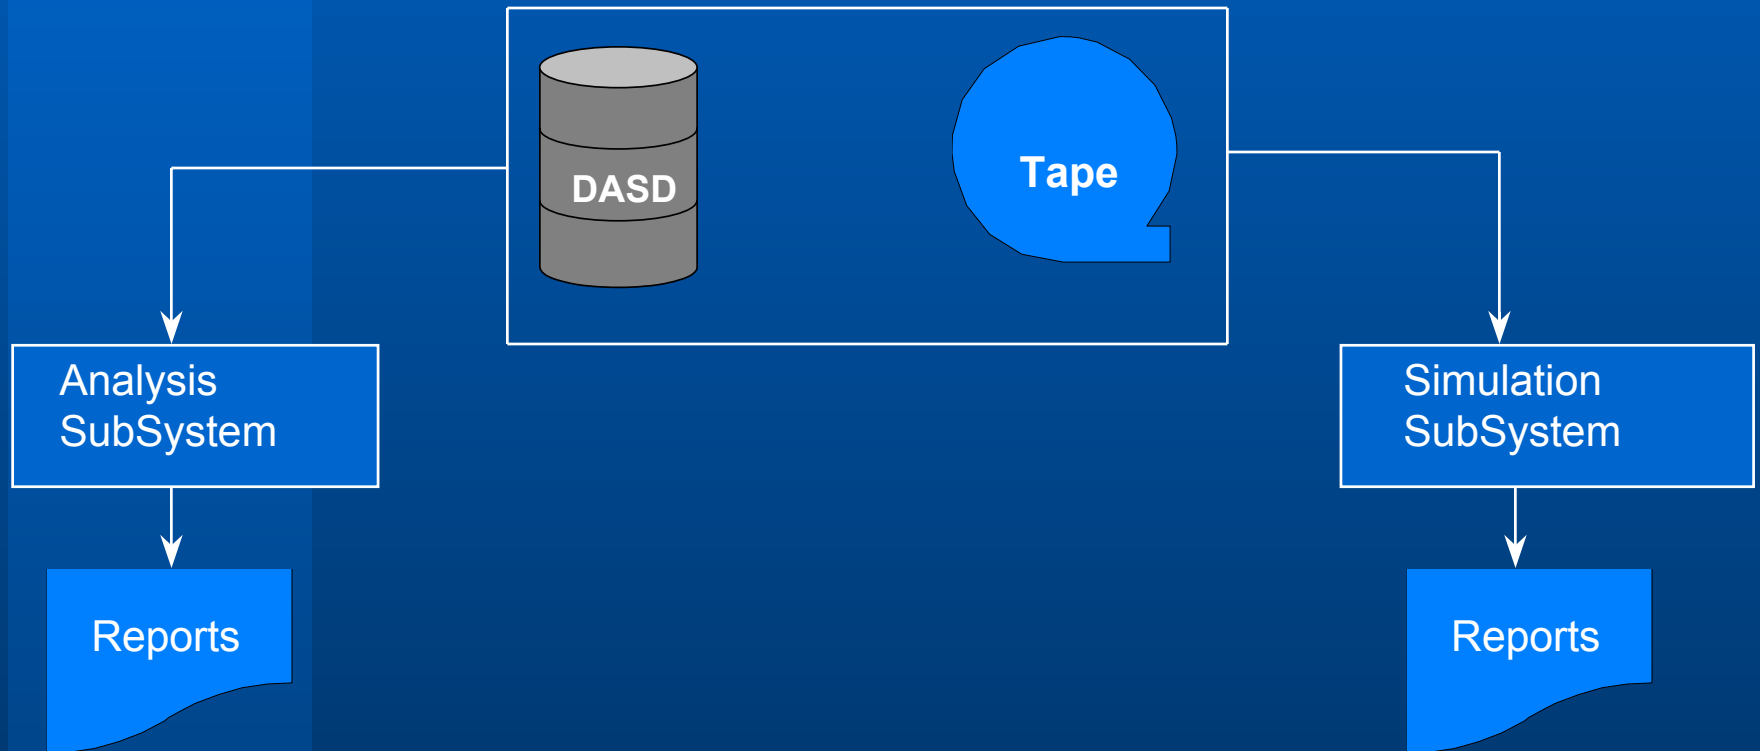
Buffer Pool Tool Components

- **Collector**
- **Statistics**
- **Simulation**
- **Utilities**
- **PC Components**
 - **BPT Graphics**
 - **Coupling Facility sizing module**

Buffer Pool Tool Components



Buffer Pool Tool Components



BufferPool Reports
Object Usage Reports

Simulation/Prediction

- BufferPool Level
- Object Level

Collector

- **Attaches to DB2 using the IFI interface, a standard supported component of DB2**
- **Starts traces, retrieves data, writes to its own highly blocked dataset**
- **Collects for a specific period, does not run all the time – yet you need volume for statistical accuracy**
 - **Small samples are not accurate (what do small and big mean?)**
 - **Combining many small samples with large time gaps violates approved statistical techniques**

Statistics

- **Shows you what happened**
 - Overall system
 - Each pool
 - Each object within every pool
- **Provides a level of information you cannot obtain from any other product**
 - Highlights application access path problems
 - Yellow Freight....
 - Other Syspacklist index problem

Simulation - *what will happen if...*

- Predicts the effect of changes

- Simple changes
 - Pool sizes
 - Pool thresholds
- Complex changes (the keys for improving performance)
 - Moving objects into different pools
 - Moving objects into a new (non-existing) pool

- Predicts

- I/O rate per second
- Hit ratios, system and application
- Object working set sizes (will be different at varying pool sizes)
- Required Pool size/Memory requirements – best performance

Simulation *Benefits*

- You don't make mistakes with your production system
 - Simple changes can be made by command and monitor the effect
 - Pool sizes
 - Pool thresholds
 - Complex changes require stop/alter/start the object, which requires an outage on that object
 - Moving objects into different pools
 - Moving objects into a new (non-existing) pool
 - **The biggest tuning improvements come from re-grouping objects into pools, based on access type & working set size**
 - Shows how objects will **interact** after changes
- Simulation has a proven track record with our clients

Utilities Provide Added Benefit

- **Application Analysis Facility**
 - Plan/Authid
- **Object Analysis**
 - Object Name
- **Connection Type Reporting**
- **Page Access Distribution Reports**

BPTGraph Speeds the Analysis

Harness the power of the workstation

- **Quick Analysis by System, Pool and Object**
 - A picture is worth a thousand words..
- **Top Ten/Worst Ten selections**
- **Cluster Analysis/Object Grouping**
- **Include/Exclude Object file generation**

BPTGraph - Overall Summary Data

Buffer Pool Tool for DB2 - BP18

Report Info | Graphic Summary | Pool Info | Object Info | Expert Tuning | Sim Graph Analysis | Sim Cluster Analysis

Collection

Date: 2003-02-13
 Time: 10:34:52
 Elapsed Time: 00:30:00

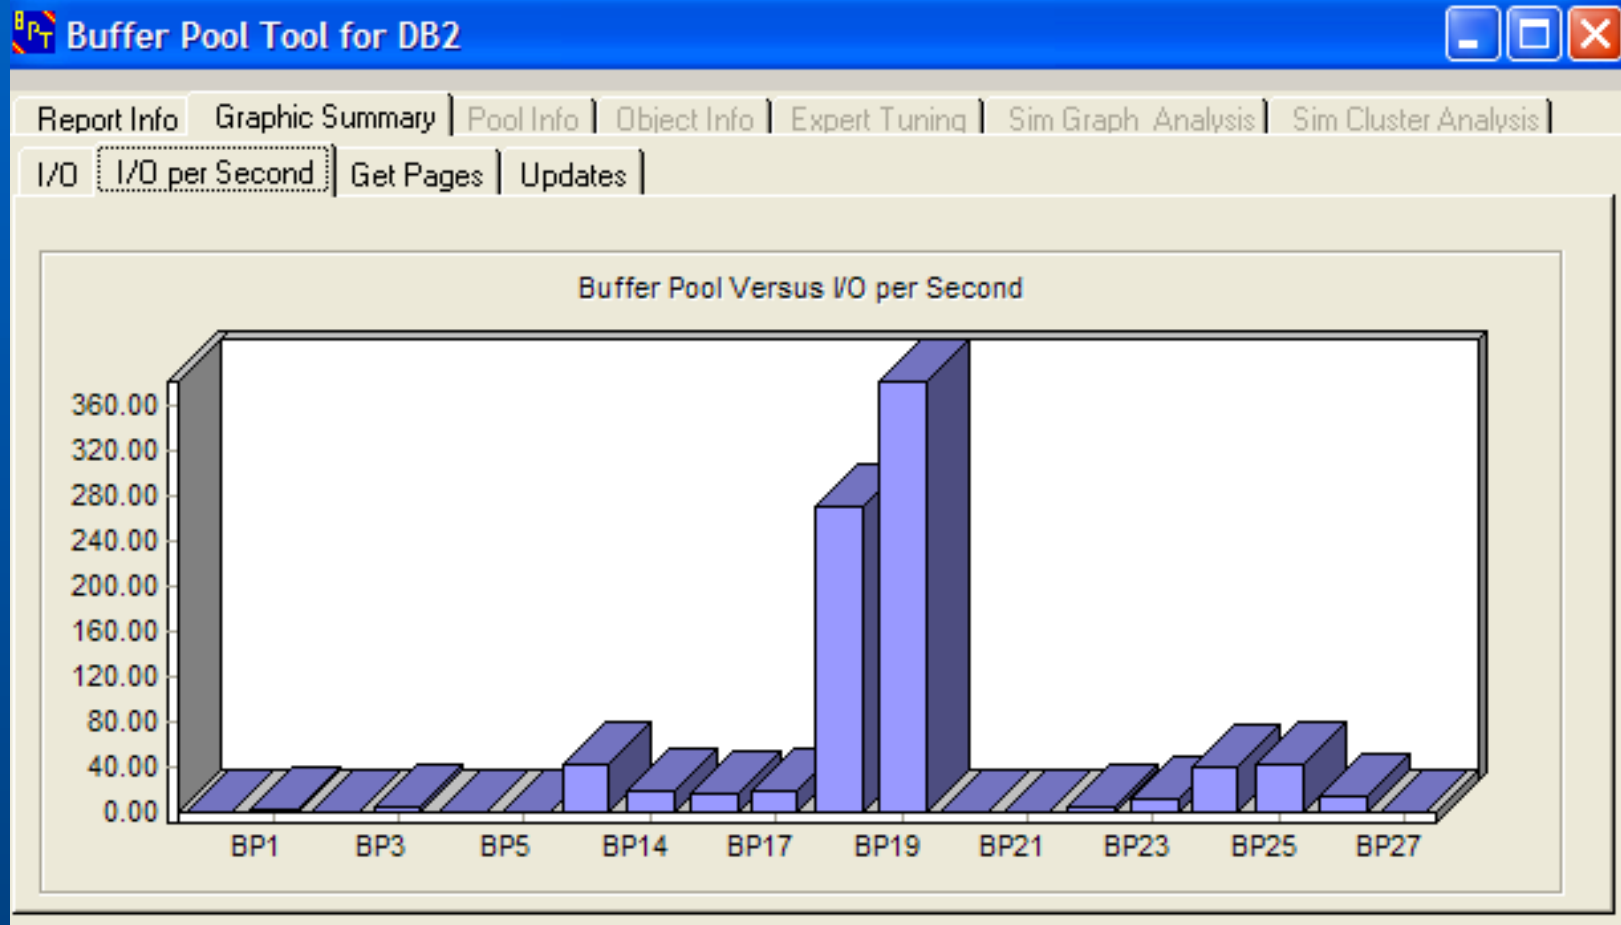
System Info

System: P20G
 Sub System: DSN1
 Db2 Version: 7.1

Pool	I/O	Get Pages	Updates	Hit Ratio	I/O Sec	Pages/Write
BP1	2791	482810	30057	99.9	1.55	5.29
BP2	377	39056	408	97.5	0.21	2.70
BP3	8362	3608218	2904659	99.3	4.65	16.70
BP4	0	386	0	100.3	0.00	0.00
BP5	0	636	0	100.2	0.00	0.00
BP6	76154	211985	85806	65.5	42.31	1.56
BP14	33562	78793	1208	58.4	18.65	1.10
BP16	30864	5937028	18232	98.6	17.15	1.58
BP17	34457	1719312	16231	97.5	19.14	1.67
BP18	488578	10213223	44031	95.3	271.43	1.22
BP19	688420	11436482	75914	90.2	382.46	1.28
BP20	7	358068	259	100	0.00	2.17
BP21	7	370000	0	100	0.00	1.22

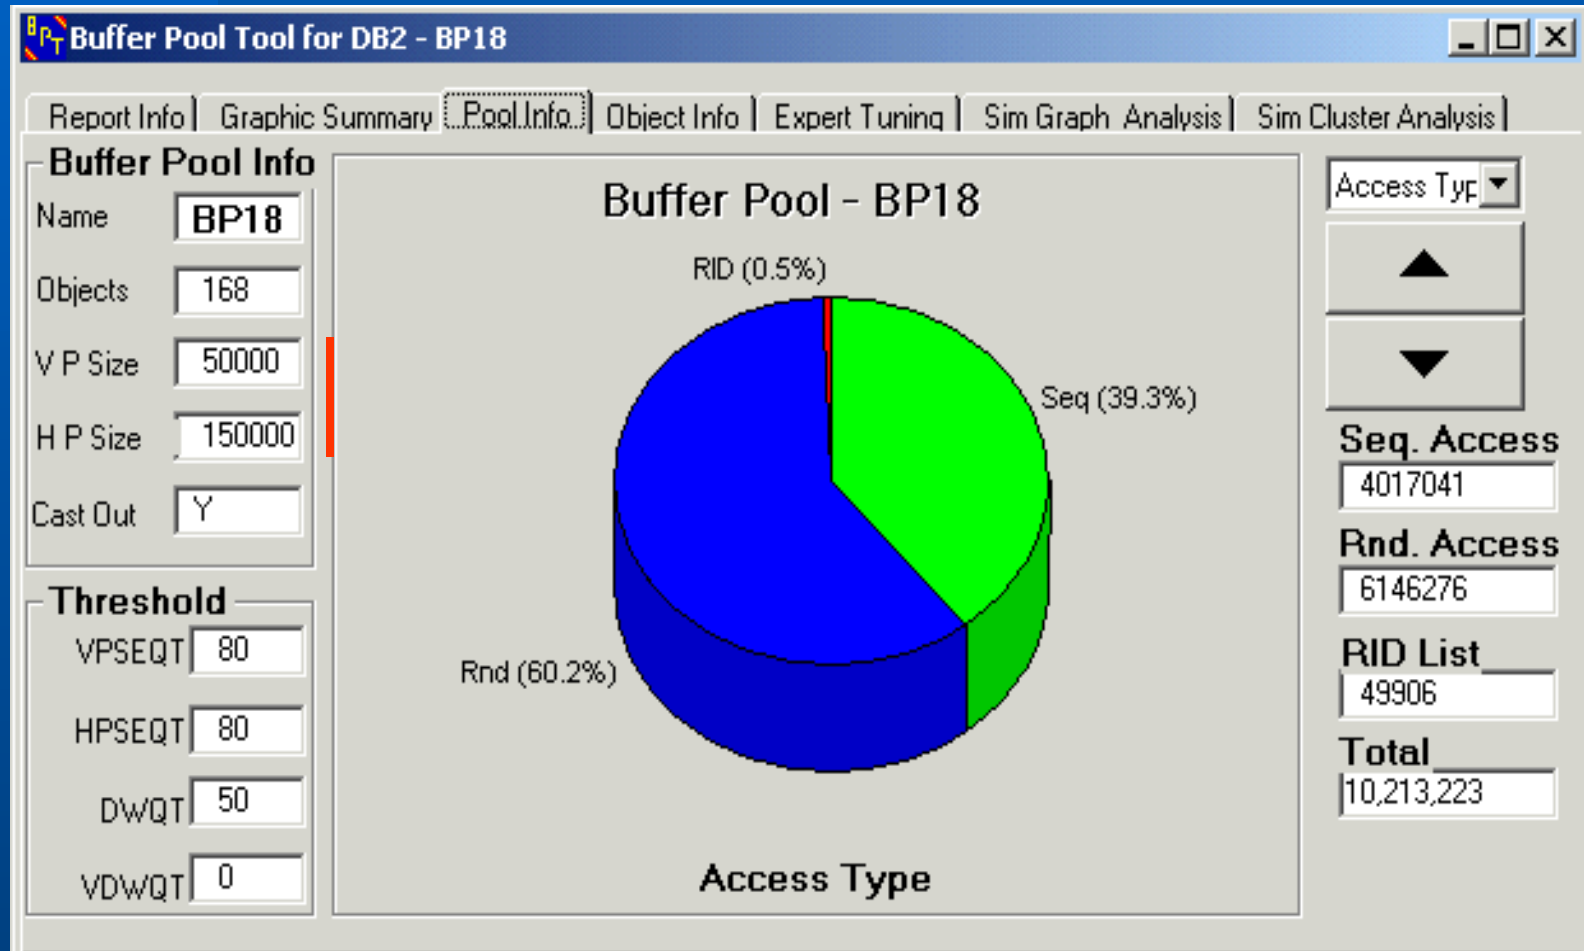
Total Read/Write IO: **1,563,837** Total Get Pages: **38,145,154**
 Overall Sys Hit Ratio: **93.55** Total I/Os per second: **868.80**
 Total Updates: **3,247,036** Pages per write: **1.91**

BPTGraph - Graphic Summary

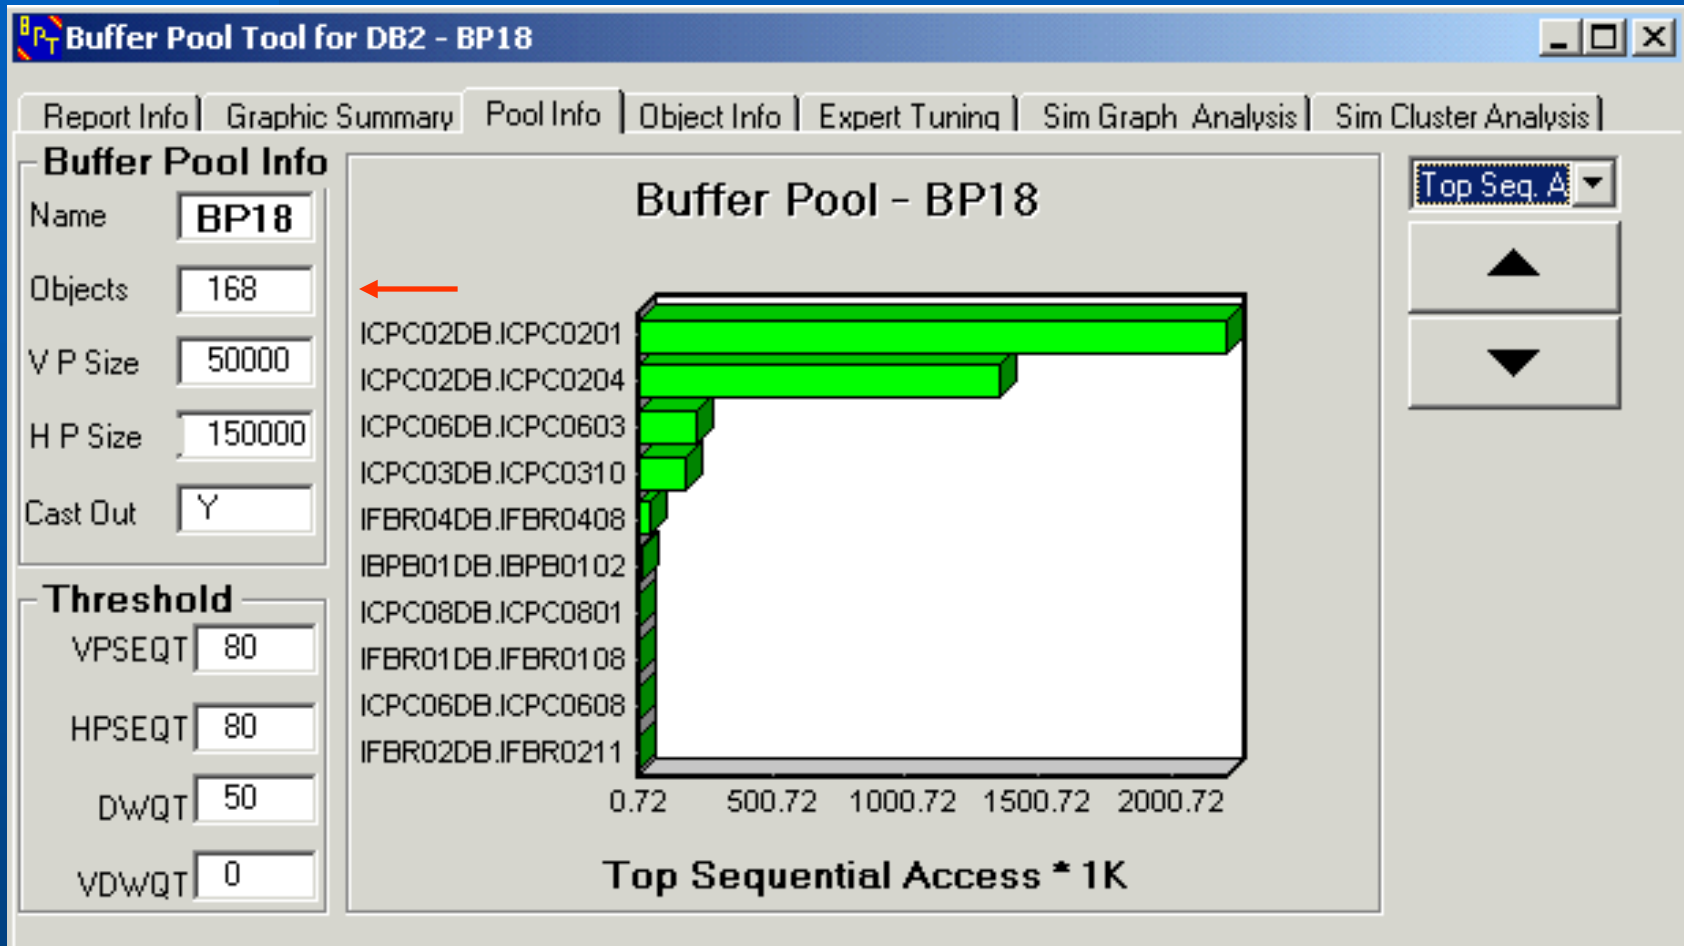


Quick graphic summary of pool performance

BPTGraph - shows pool usage



BPTGraph - shows sequential impactors



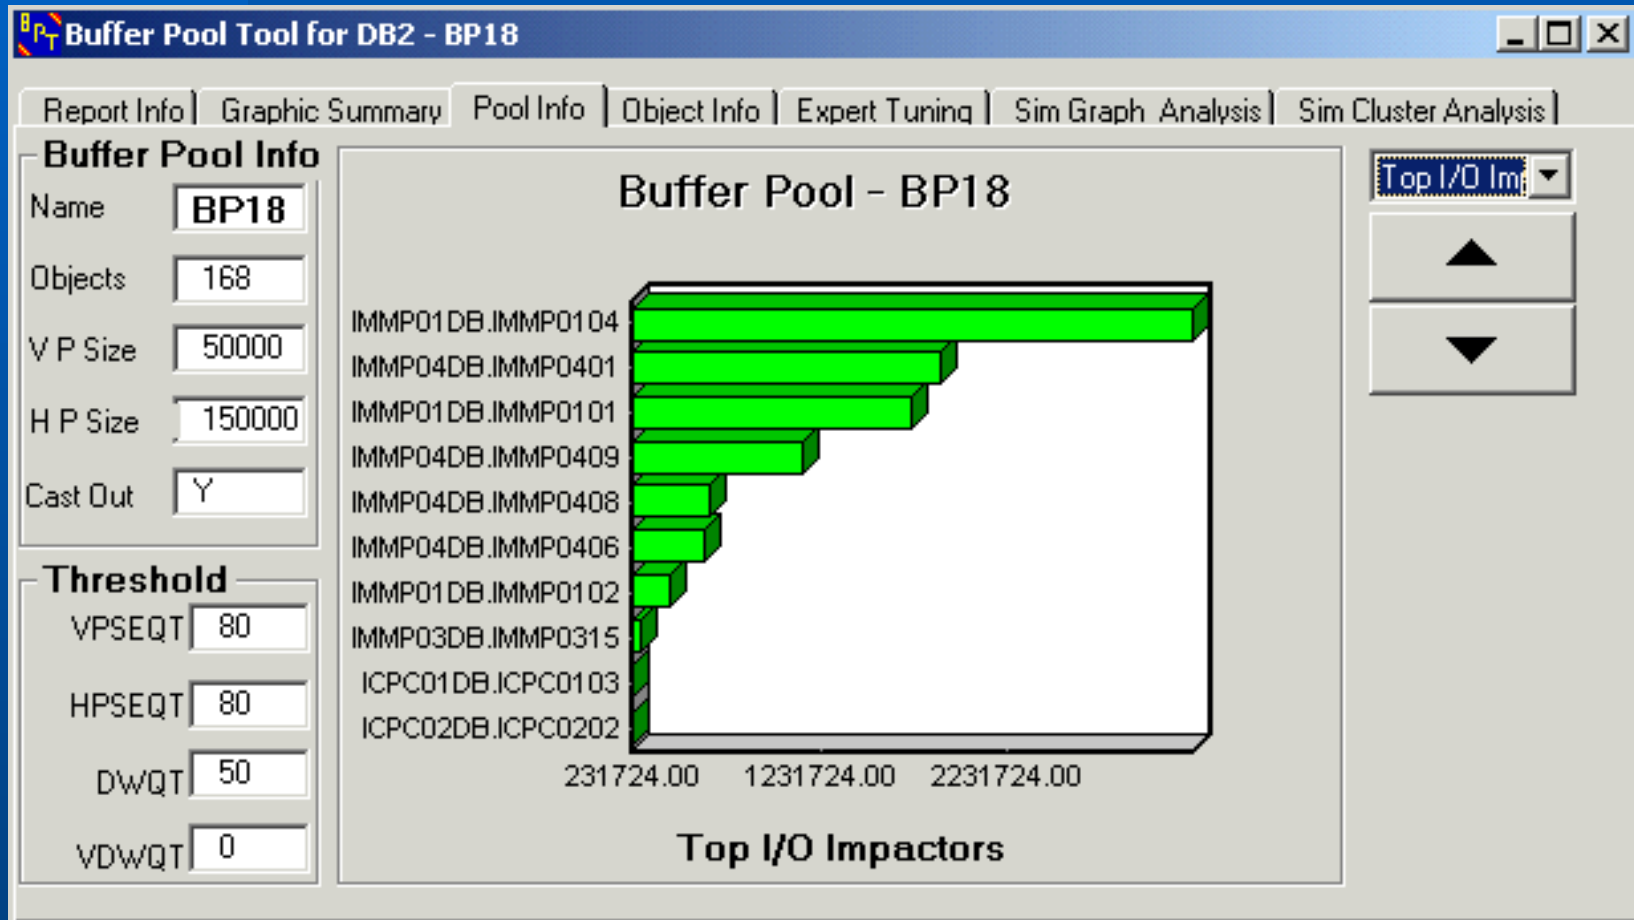
BPTGraph - shows object detail

The screenshot displays the 'Buffer Pool Tool for DB2 - BP18' window. The 'Object Info' tab is selected, showing various performance metrics for object BP18. The interface is organized into several sections:

- Buffer Pool Info:** Name: BP18, Objects: 168, V P Size: 50000, H P Size: 150000, Cast Out: Y.
- Threshold:** VPSEQT: 80, HPSEQT: 80, DWQT: 50, VDWQT: 0.
- Performance Metrics:**
 - Total Get Pages: 2264240
 - Pages Read Sync: 1
 - App Hit Ratio: 100
 - Get Page Rand: 55750
 - Pages Read Seqpr: 0
 - System Hit Ratio: 100
 - Get Page Seq: 2208490 (indicated by a red arrow)
 - Pages Read Listpr: 0
 - Read IO Rate/sec: 0.00 (indicated by a red arrow)
 - Get Page RidList: 0
 - Pages Read Dynpr: 0
 - Pages / Write: 1.08
 - Avg Synchron IO (ms): 48.00 (indicated by a red arrow)
 - Avg SP IO (Seq Pref): 0.00
- Buttons:** A 'Close' button is located at the bottom right.

This object lives in the pool


BPTGraph - shows I/O impactors



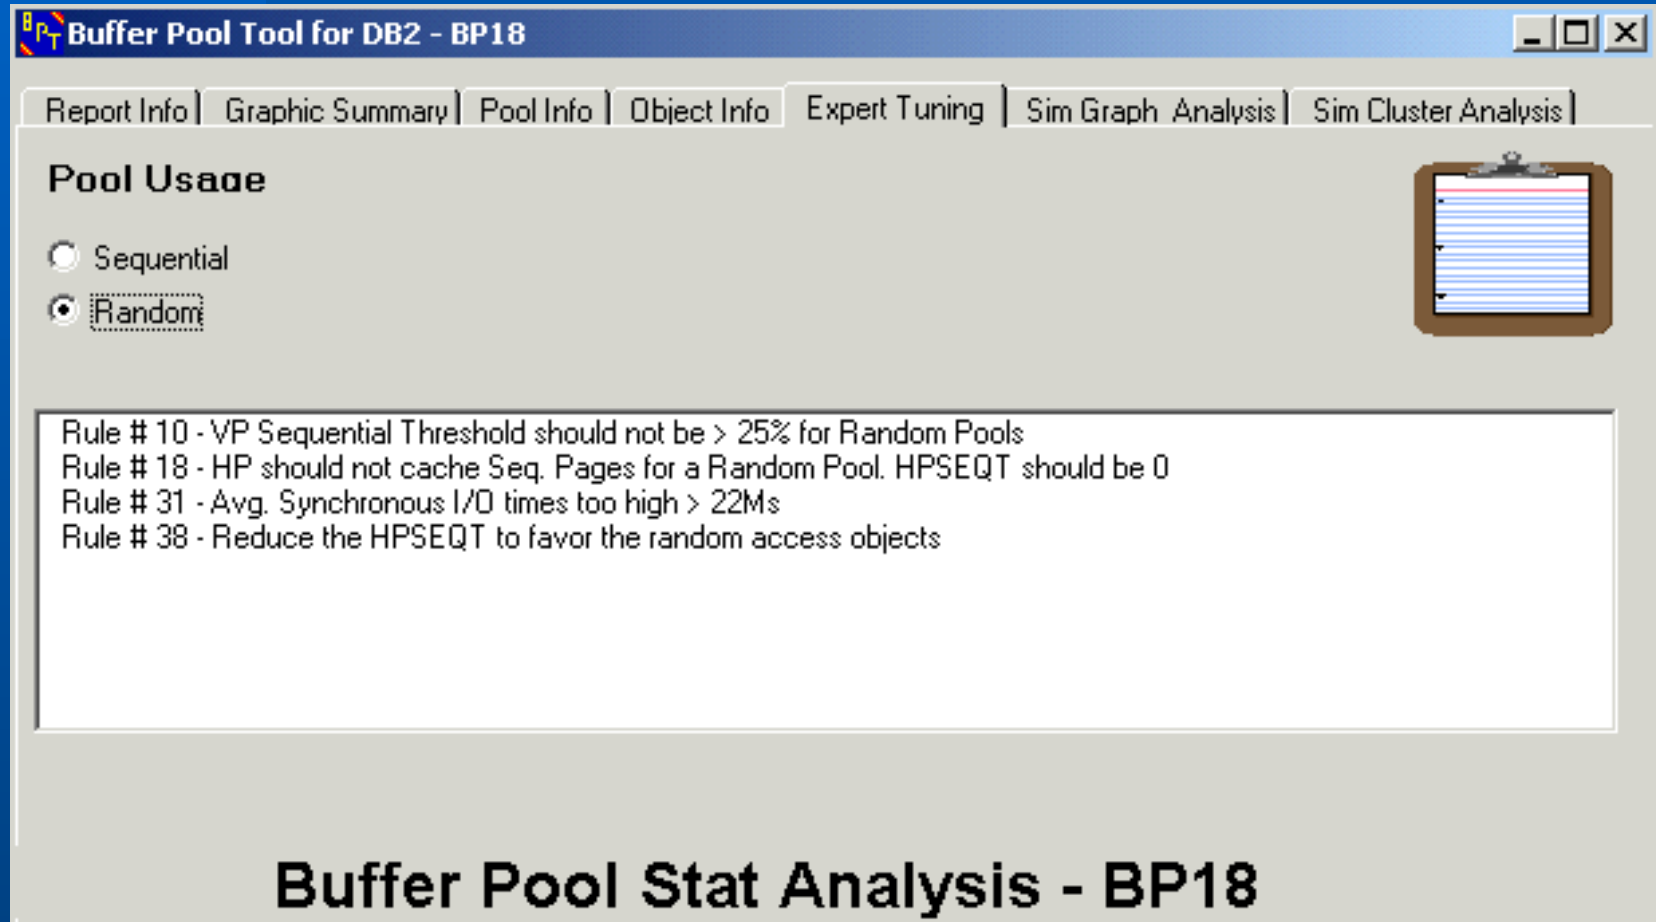
Not the same as the I/O Rate

BPTGraph - Object Detail Statistics

Click & Sort



BPTGraph - Expert Tuning Suggestions




The screenshot shows a software window titled "Buffer Pool Tool for DB2 - BP18". The window has a menu bar with the following items: "Report Info", "Graphic Summary", "Pool Info", "Object Info", "Expert Tuning", "Sim Graph Analysis", and "Sim Cluster Analysis". The "Expert Tuning" menu item is currently selected. Below the menu bar, the text "Pool Usage" is displayed. There are two radio button options: "Sequential" and "Random". The "Random" option is selected. To the right of these options is a clipboard icon. Below the radio buttons, there is a text box containing the following expert tuning suggestions:

- Rule # 10 - VP Sequential Threshold should not be > 25% for Random Pools
- Rule # 18 - HP should not cache Seq. Pages for a Random Pool. HPSEQT should be 0
- Rule # 31 - Avg. Synchronous I/O times too high > 22Ms
- Rule # 38 - Reduce the HPSEQT to favor the random access objects


At the bottom of the window, the text "Buffer Pool Stat Analysis - BP18" is displayed.

BPTGraph - shows I/O reduction/payback



50,000 buffers saves 10 I/O a Sec

BPTGraph - Hit Ratio is not a useful metric



BPTGraph - shows you the problem objects

The screenshot shows the 'Buffer Pool Tool for DB2 - BP18' window. The 'Pool Usage Intent' is set to 'Sequential'. The 'Pool Size' is 200000 and the 'Cluster Radius' is 1.2. Under 'Cluster Info', there are two tables: 'Clusters' and 'Objects'.

Pool Usage Intent
 Sequential Random

Pool Size: 200000
Cluster Radius: 1.2

Cluster Info

Clusters

Object	Smallest Ma	Largest Max
1	16547	16547
2	3615	4415
3	1	1122

Objects

Type	Object	Max Work Set
T	ICPC02DB.ICPC0204	16547

BPTGraph -

Buffer Pool Tool for DB2 - BP0



Report Info | Graphic Summary | Pool Info | Object Info | Expert Tuning |

Collection

Date 2004-03-01

Time 10:05:30

Elapsed Time 00:15:48

System Info

System NENT

Sub System NBP


DB2 Version 7.1

Pool	I/O	Get Pages	Updates	Hit Ratio	RI/O/Sec	WIO/Sec
BP0	86914	792976	60	77	91.65	0.
BP2	2989001	20443958	76353	-18.6	3,122.22	30.
BP3	1304487	16649309	112401	59.4	1,321.93	54.
BP4	836	39610	9288	96.4	0.81	0.
BP6	619697	1602799	7192	-63.1	651.08	2.
BP7	41327	262998	23924	68.4	37.85	5.
BP13	39	222	104	88.3	0.03	0.
BP30	5701	1245951	213501	98	3.45	2.
BP49	0	500	300	100.2	0.00	0.
BP32K	10560	82445	19188	84.7	9.29	1.

Total Read/Write IO	5,058,562	Total Get Pages	41,120,768
Overall Sys Hit Ratio	17.50	Total I/Os per second	5,336.04
Total Updates	462,311	Pages per write	2.14

Huge I/O rate

BPTGraph -



Pool too small

BPTGraph -

Buffer Pool Tool for DB2 - BP0



Report Info | Graphic Summary | Pool Info | Object Info | Expert Tuning |


Buffer Pool Info

Name
Objects
VP Size
HP Size
Cast Out

Threshold

VPSEQT
HPSEQT
DWQT
VDWQT

Buffer Pool - BP0



Top I/O Ra



Pool too small

BPTGraph -

BPT Buffer Pool Tool for DB2 - BP0

Report Info | Graphic Summary | Pool Info | Object Info | Expert Tuning

Buffer Pool Info

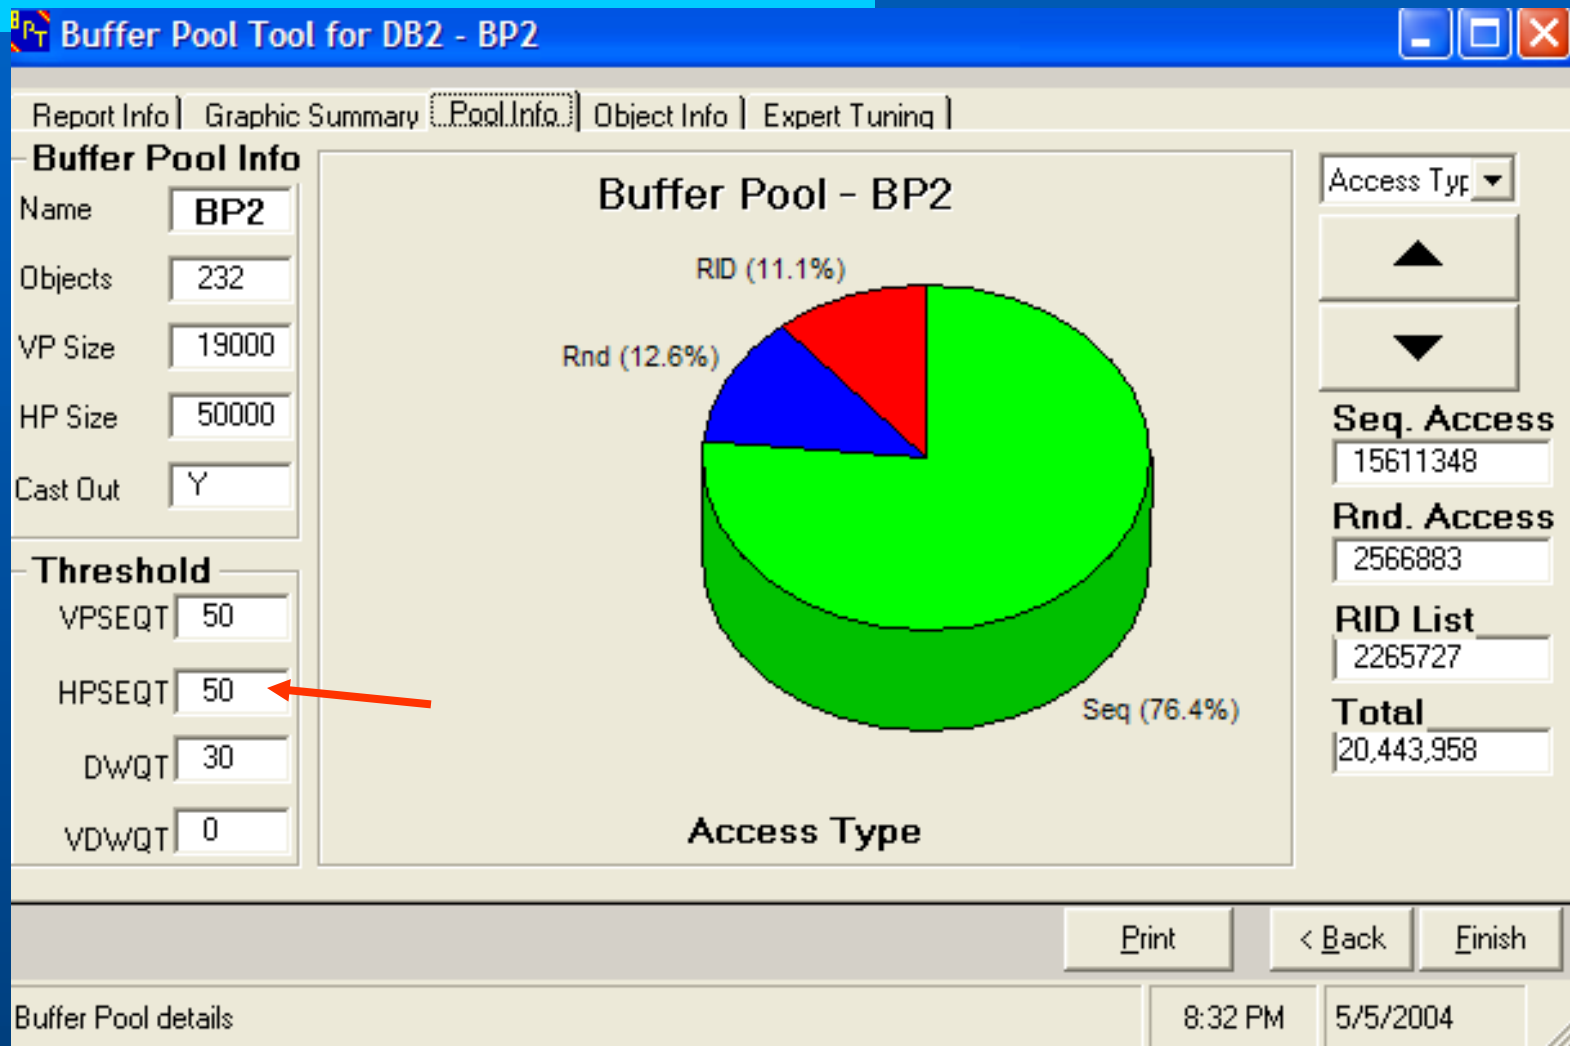
Name	BP0
Objects	34
VP Size	2000
HP Size	0
Cast Out	Y

Threshold

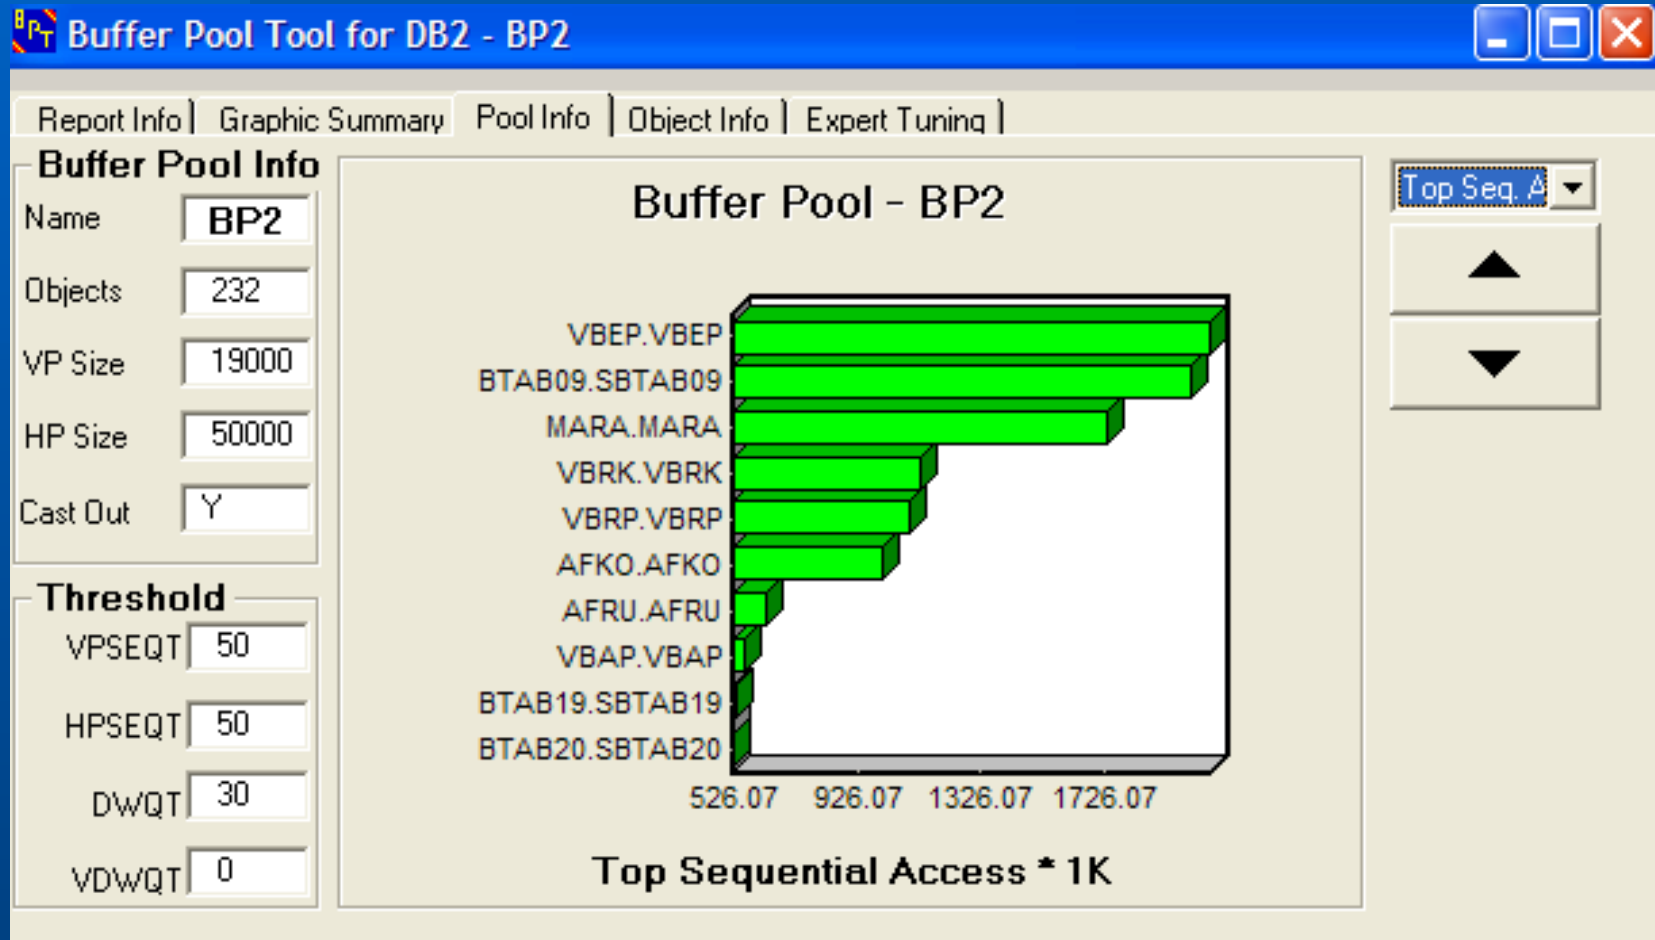
VPSEQT	50
HPSEQT	80
DWQT	50
VDWQT	10

Total Get Pages	518890	Pages Read Sync	73225	App Hit Ratio	85.9
Get Page Rand	421610	Pages Read Seqpr	91761	System Hit Ratio	67.9
Get Page Seq	94600	Pages Read Listpr	0	Read IO Rate/sec	80.60
Get Page RidList	2680	Pages Read Dynpr	1718	Pages / Write	0.00
Avg Synch IO (ms)	2.00	Avg SP IO (Seq Pref)	4.00	Close	

BPTGraph -



BPTGraph -



Pool too small, maybe other problems?

BPTGraph -

Buffer Pool Tool for DB2 - BP2

Report Info | Graphic Summary | **Pool Info** | Object Info | Expert Tuning

Buffer Pool Info

Name: **BP2**

Objects: **232**

VP Size: **19000**

HP Size: **50000**

Cast Out: **Y**

Threshold

VPSEQT: **50**

HPSEQT: **50**

DWQNT: **30**

VDWQNT: **0**

Total Get Pages 2094085	Pages Read Sync 15199	App Hit Ratio 99.3
Get Page Rand 25199	Pages Read Seqpr 2089942	System Hit Ratio -3.3
Get Page Seq 2068886	Pages Read Listpr 0	Read IO Rate/sec 86.93
Get Page RidList 0	Pages Read Dynpr 58461	Pages / Write 1.61
Avg Synchron IO (ms) 2.00	Avg SP IO (Seq Pref) 4.00	Close

BPTGraph -

Buffer Pool Tool for DB2 - BP2

Report Info | Graphic Summary | **Pool Info** | Object Info | Expert Tuning

Buffer Pool Info

Name	BP2
Objects	232
VP Size	19000
HP Size	50000
Cast Out	Y

Threshold

VPSEQT	50
HPSEQT	50
DWQT	30
VDWQT	0

Total Get Pages	1796540	Pages Read Sync	73449	App Hit Ratio	95.9
Get Page Band	30231	Pages Read Seqpr	1606169	System Hit Ratio	3.5
Get Page Seq	1732519	Pages Read Listpr	32738	Read IO Rate/sec	133.57
Get Page RidList	33790	Pages Read Dynpr	22024	Pages / Write	1.00
Avg Synch IO (ms)	2.00	Avg SP IO (Seq Pref)	4.00	<input type="button" value="Close"/>	

BPTGraph -

Buffer Pool Tool for DB2 - BP2



Report Info | Graphic Summary | Pool Info | Object Info | Expert Tuning |

Buffer Pool Info

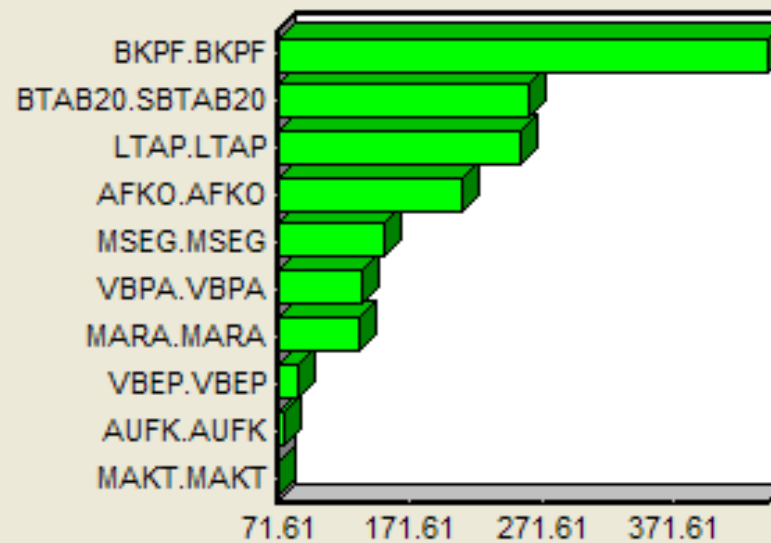
Name
Objects
VP Size
HP Size
Cast Out

Threshold

VPSEQT
HPSEQT
DWQT
VDWQT

Buffer Pool - BP2

Top I/O Re



Highest IO Rates/Sec

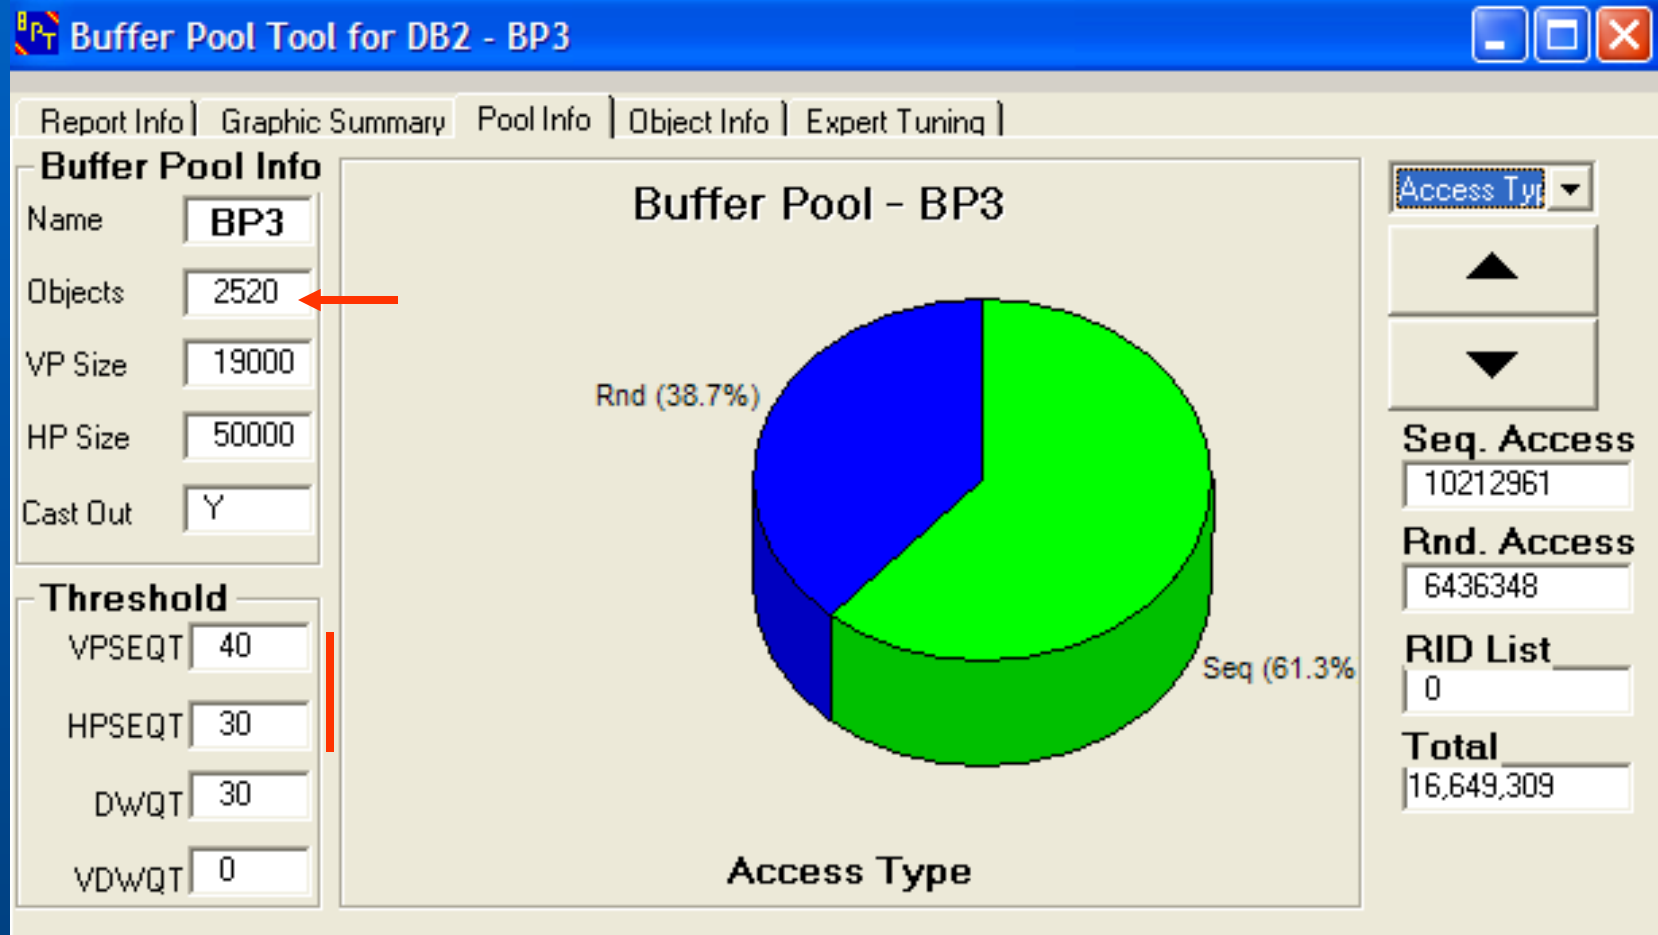
BPTGraph - object BKPF

The screenshot shows the 'Buffer Pool Tool for DB2 - BP2' application window. The 'Object Info' tab is selected, displaying various performance metrics for buffer pool BP2. The metrics are organized into several columns and rows. Two red arrows highlight specific values: one points to 'Pages Read Sync' (401151) and another points to 'Read IO Rate/sec' (443.54).

Section	Metric	Value
Buffer Pool Info	Name	BP2
	Objects	232
	VP Size	19000
	HP Size	50000
	Cast Out	Y
	Threshold	VPSEQT
HPSEQT		50
DWQT		30
VDWQT		0
Total Get Pages: 544113		
Pages Read Sync: 401151		
App Hit Ratio: 26.3		
Get Page Band: 29818		
Pages Read Seqpr: 517170		
System Hit Ratio: -84.7		
Get Page Seq: 514295		
Pages Read Listpr: 0		
Read IO Rate/sec: 443.54		
Get Page RidList: 0		
Pages Read Dynpr: 86860		
Pages / Write: 1.17		
Avg Synch IO (ms): 2.00		
Avg SP IO (Seq Pref): 5.00		
Close		

Synch I/O is 13x Random Access³⁶

BPTGraph - similar pool usage problem




BPTGraph - Another system

Pool	I/O	Get Pages	Updates	Hit Ratio	RIO/Sec	WIO/Sec
BP0	367	4099581	711	100	0.50	0.
BP1	1884504	23941040	1632489	89.8	2,733.72	407.
BP2	906410	7219383	1208322	61.9	1,437.72	72.
BP3	399	508375	286517	100	0.00	0.
BP4	351021	419747	227	-20.6	584.76	0.
BP5	220	132003	0	99.8	0.37	0.
BP6	187	887	137	83.3	0.25	0.
BP7	12325	99926	3447	75.5	19.76	0.
BP8	206677	1203809	119431	83	282.33	62.
BP9	67943	675981	46	89.9	113.24	0.
BP10	64820	86010	85	24.6	108.03	0.
BP11	92673	113915	52606	23.9	76.73	77.

Total Read/Write IO	3,587,546	Total Get Pages	38,500,657
Overall Sys Hit Ratio	84.05	Total I/Os per second	5,979.24
Total Updates	3,304,018	Pages per write	2.06

512,000 buffers overall, > 2 Gig of memory

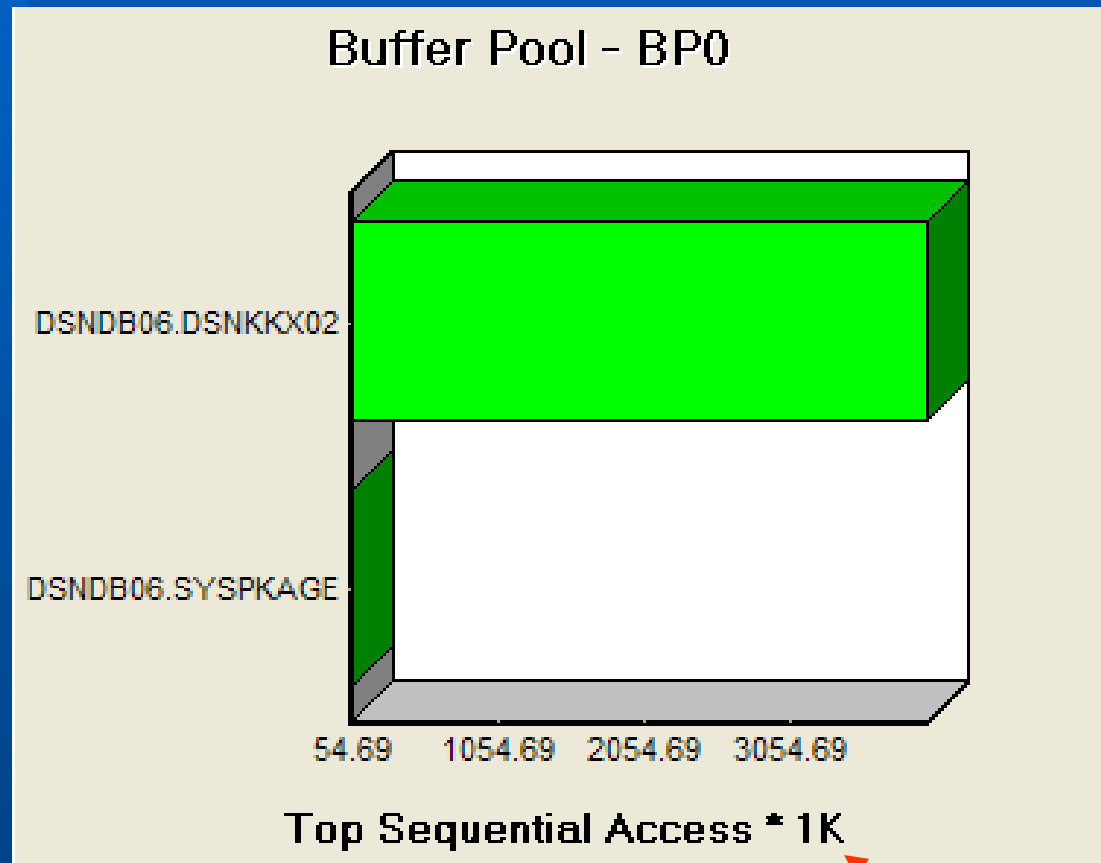
BPTGraph - What's happening here?



Surprise, surprise

4 million sequential access?
> 10% of the whole system???

BPTGraph - Index on SYSPACKAGE




BPTGraph - Memory resident

Total Get Pages 3991395	Pages Read Sync 0	App Hit Ratio 100
Get Page Rand 0	Pages Read Seqpr 0	System Hit Ratio 100
Get Page Seq 3991395	Pages Read Listpr 0	Read IO Rate/sec 0.00
Get Page RidList 0	Pages Read Dynpr 0	Pages / Write 0.00
Avg Synch IO (ms) 0.00	Avg SP IO (Seq Pref) 0.00	<input type="button" value="Close"/>


Costing 6-8 Mips

BPTGraph - Big pool, dataspace



All Indexes

BPTGraph - Reasonable payback ?



5% reduction

50K buffers, saves 150 I/O sec

BPTGraph - Another system..

BPT Buffer Pool Tool for DB2

Report Info | Graphic Summary | Pool Info | Object Info | Expert Tuning | Sim Graph Analysis | Sim Cluster Analysis

Collection

Date: 2003-04-25
Time: 00:30:34
Elapsed Time: 00:59:54

System Info


System: APB2
Sub System: DPR1
DB2 Version: 7.1

Pool	I/O	Get Pages	Updates	Hit Ratio	RI0/Sec	WIO/Sec
BP0	5295	625794	2344	97.3	1.34	0.
BP1	129690	521373	116586	-76.2	30.69	5.
BP2	148	4823287	3273660	100	0.04	0.
BP3	549991	10228436	306689	44.4	136.81	16.
BP4	734741	18943183	401976	76.5	182.32	22.
BP5	38213	1368285	13645	83.9	8.70	1.
BP6	275546	2748354	67464	42.3	70.73	5.
BP7	179277	4386248	29213	5.7	48.56	1.
BP8	107	4178	31	85.6	0.03	0.
BP10	129449	969665	55189	-40.1	32.72	3.
BP32K	4	31	16	103.2	0.00	0.

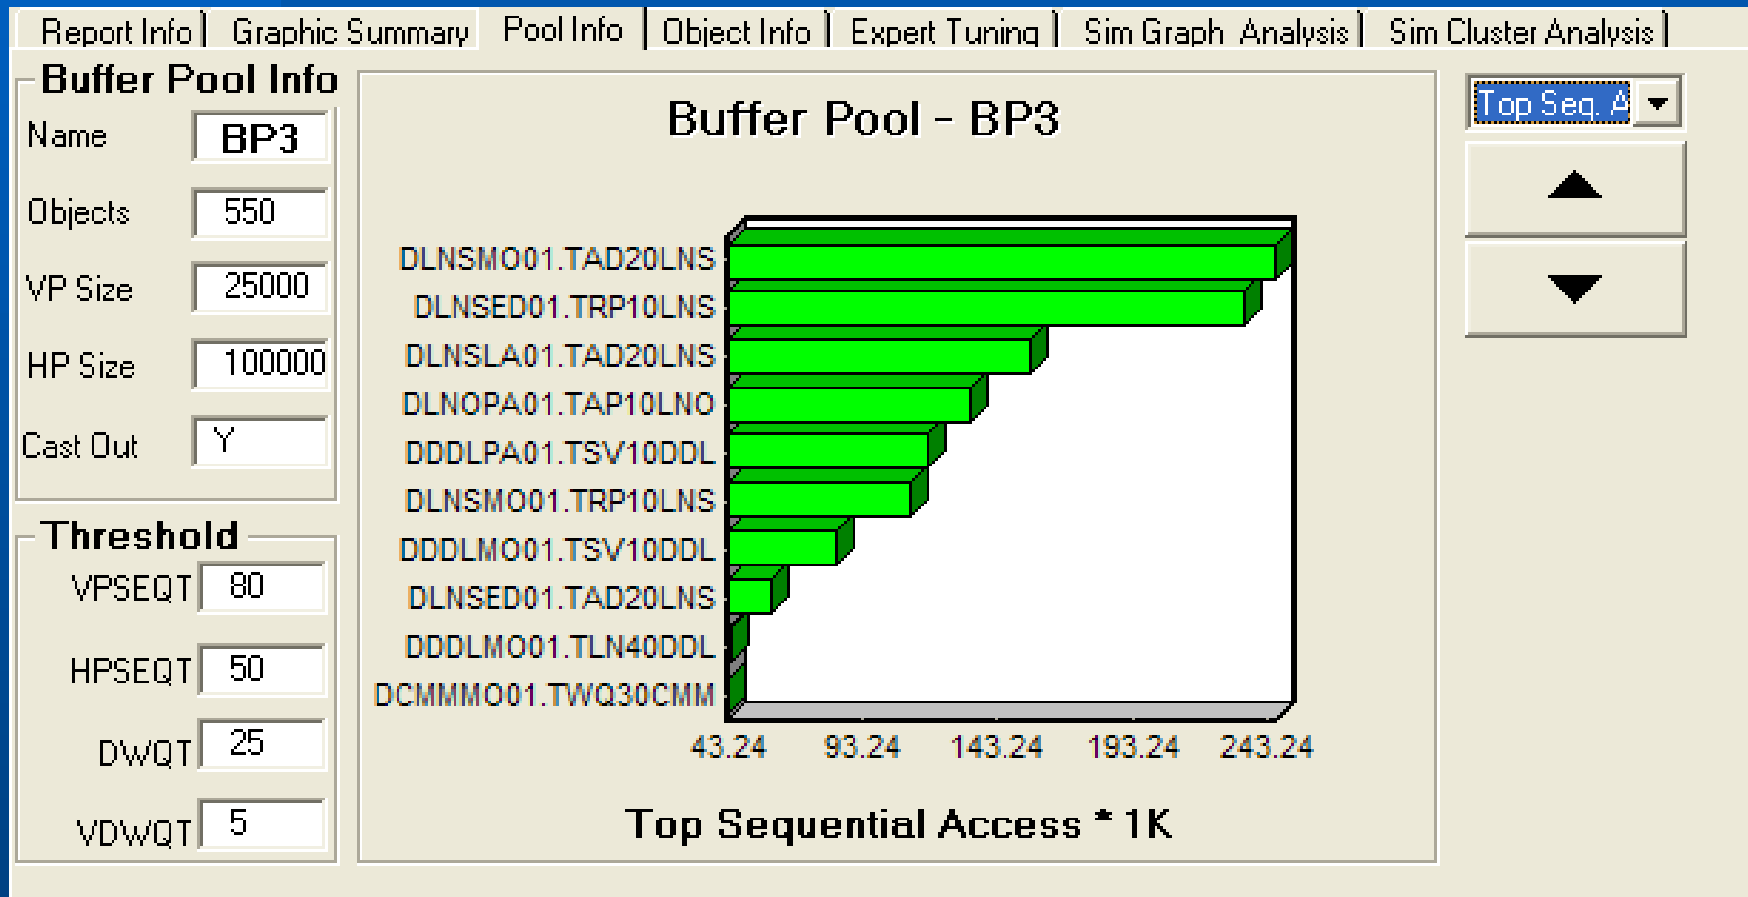
< ||| >

Total Read/Write IO	2,042,461	Total Get Pages	44,618,834
Overall Sys Hit Ratio	58.83	Total I/Os per second	568.30
Total Updates	4,266,813	Pages per write	3.16


BPTGraph -.. Where's the problem?



BPTGraph -.. Where's the problem?



BPTGraph - Some small gain



BPTGraph - Last system example

Buffer Pool Tool for DB2 - BP2

Report Info | Graphic Summary | Pool Info | Object Info | Expert Tuning | Sim Graph Analysis | Sim Cluster Analysis

Collection

Date: 2003-03-28
 Time: 15:22:06
 Elapsed Time: 00:58:09


System Info

System: SYS1
 Sub System: DB01
 DB2 Version: 7.1

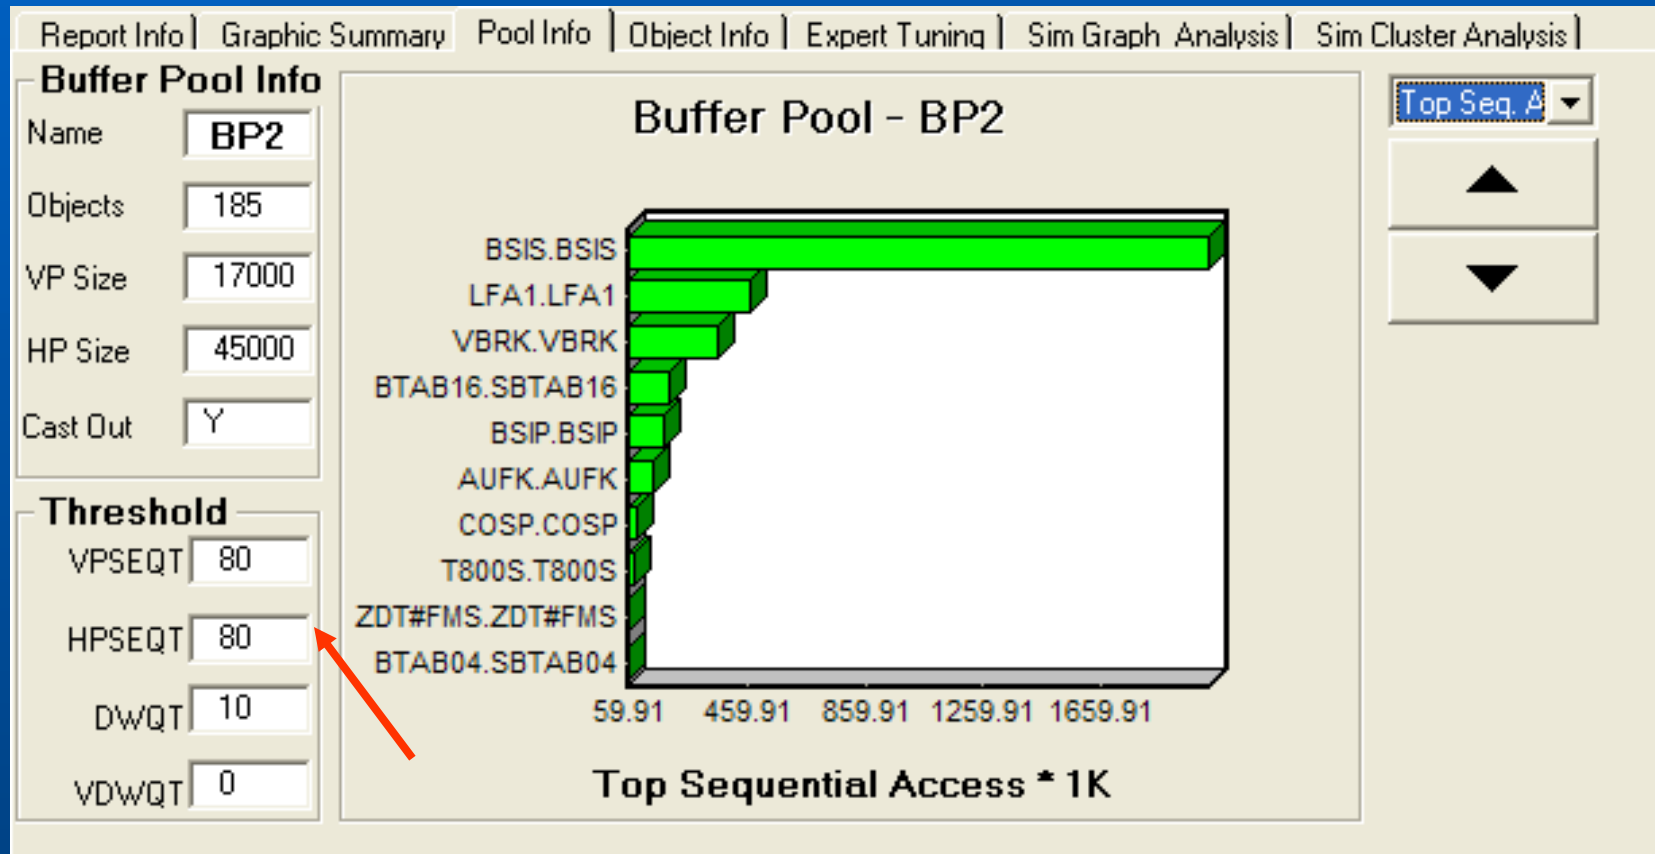
Pool	I/O	Get Pages	Updates	Hit Ratio	RIO/Sec	WIO/Sec	Pages/Write
BP0	627	8085	215	92.8	0.16	0.02	2.24
BP1	2498	285414	246080	98.7	0.18	0.54	15.84
BP2	910275	6305745	165309	35.7	248.05	12.85	1.86
BP3	342429	12283232	710194	85	66.00	32.14	2.57
BP4	69097	1235173	305302	96.7	10.47	9.34	2.03
BP5	952	180295	93559	99.5	0.24	0.04	12.72
BP6	80393	264913	6663	20.7	22.89	0.15	5.28
BP7	30405	175029	16726	32.3	8.17	0.55	2.79
BP10	30909	542523	14871	83.9	7.66	1.20	2.12
BP11	48690	124815	4987	-38.4	13.88	0.07	5.70
BP12	4083	409667	18273	99	0.95	0.22	5.14
BP14	81848	217214	4486	-76.9	23.37	0.09	4.47
BP15	6006	3152388	0	99.7	1.72	0.00	0.00

Total Read/Write IO	1,623,452	Total Get Pages	25,436,064
Overall Sys Hit Ratio	72.64	Total I/Os per second	465.31
Total Updates	1,595,929	Pages per write	2.47


BPTGraph - Last system example



BPTGraph - Last system example



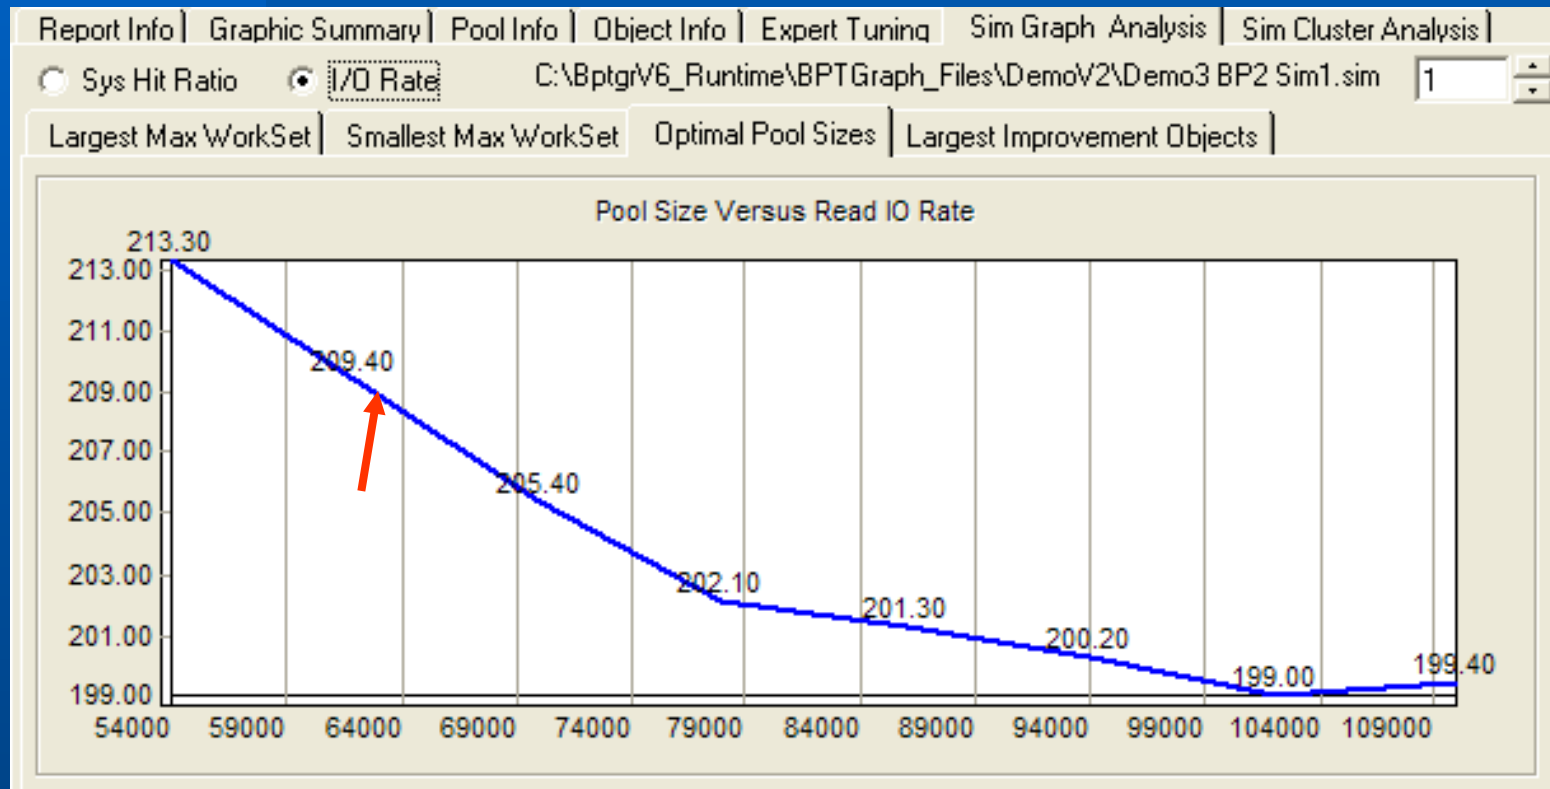
BPTGraph - Just add some memory?



Memory will help....

BPTGraph - Better Performance

Same or less memory....



Moving the objects provides a much larger gain

BPTGraph - Better Performance

By moving the 2 large SPs

Report Info | Graphic Summary | Pool Info | Object Info | Expert Tuning | Sim Graph Analysis | Sim Cluster Analysis

Pool Usage Intent **Pool Size** **Cluster Radius**

Sequential Random

C:\Bptq\V6 Runtime\BPTGraph Files\Demo\2\Demo3 BP2 Exckude.sim

Cluster Info

Clusters			Objects		
Object	Smallest Ma	Largest Max	Type	Object	Max Work Set
1	36905	47949	T	BSIS.BSIS	47949
2	5248	13055	T	CO SP.CO SP	36905
3	10	3327			

BPTGraph - shows objects that belong together

The proper grouping of objects into pools is the proven approach for optimizing the memory/performance trade-off

Prediction/Simulation shows the benefits of proper grouping, both for I/O reduction, and for Hit Ratio

Cluster Analysis shows you which objects belong together, and which should be moved out of a pool

Buffer Pool

Tuning Steps

- **Select your busy periods**
 - Monitor performance, keep performance history
 - Both system and application
- **Collect performance data**
- **Evaluate statistics, use BPTGraph to highlight the pools with highest I/O rates**
- **Identify the objects causing the most I/Os**
 - Simulate pool size changes, use BPTGraph to evaluate both size changes, and moving objects to different pools
 - Implement changes that provide the greatest benefit
 - Monitor performance, compare to original system & application performance

Buffer Pool Tool *Benefits*

- **Reduced transaction response times**
 - Average of 15% across the client base
- **Reduced batch elapsed times**
 - In some cases, more than an hour
- **Improved client and user productivity**
- **Saves CPU \$\$ through I/O elimination**
- **Better usage of memory resources**
- **Highlights application problems not shown by other products**

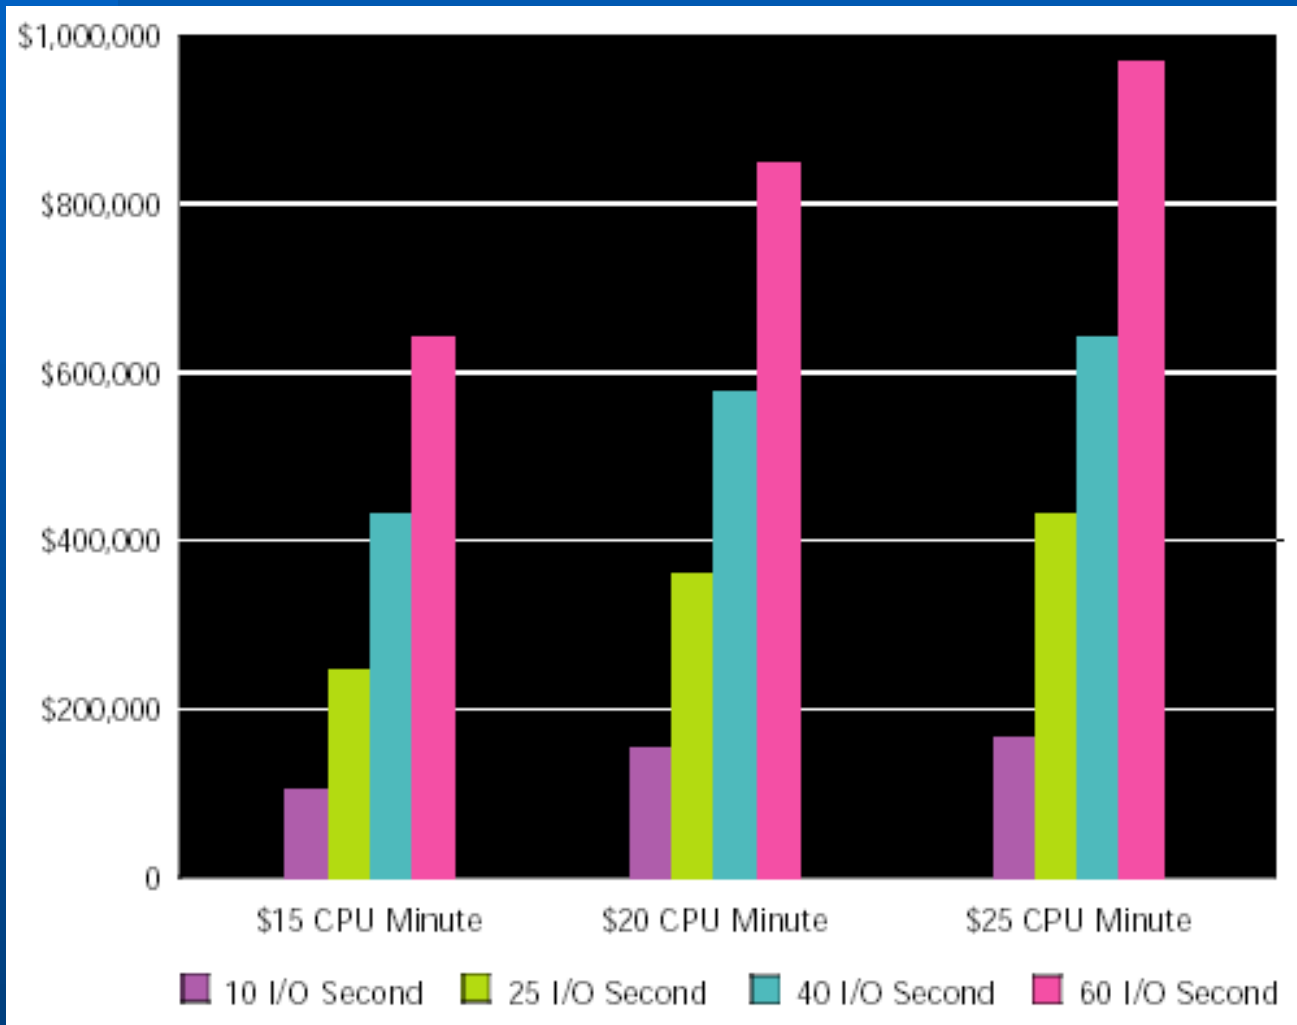
Buffer Pool Tool *Benefits*

- Provides the only proven methodology for pool tuning
- Removes the guesswork
- Helps educates your staff
- Helps leverage your people resources
 - Time is money

Buffer Pool Tool will Help You...

- **Leverage your costly Hardware investments**
- **Optimize your performance from available resources**
- **Improve service to your internal and external clients**
- **Reduced online response times provide**
 - **Improved productivity – cost savings**
 - **Clients get a better impression of the corporation**
 - **A better impression provides more long term growth and profitability**

CPU Saving from I/O Elimination



Buffer Pool Tool

- Does not try to be everything to everybody
- Helps you learn about your system & applications
- Helps you learn about DB2 performance
- Expert tuning is based upon accepted thresholds, and the performance of your system
- Shows you the underlying reasons for things, and **helps you make the best performance choices**

Coupling Facility Sizing Module

- **Imports Statistics from multiple systems**
- **Provides the Initial Sizings for all CF Structures**
- **New sizings when planning to add another member**

Included component of Buffer Pool Tool

Buffer Pool Optimizer

- Automates the simulation/prediction process
- Optimize performance for pools or objects
- Generates the statements to move objects, set thresholds, and number of buffers
- Optimizes for three performance periods
- Provides sets of control statements to change pool sizes and thresholds for different performance periods
 - Online vs batch

Summary

- Proven Solution, the industry standard for years
- Real client base, with client references *you can actually speak with...*
- Clients with measured performance improvements
- Clients that will talk about their cost savings
- Running on DB2 V8 right now, we have worked on every version prior to GA.

Caution for the future....

- **Don't be fooled and think that the need for tuning will go away with 64-bit processors and large memory**
- **Throwing memory at DB2 and pools does not get rid of performance problems**

Questions?

What additional information can I provide?

Is there anything else you would like to see?