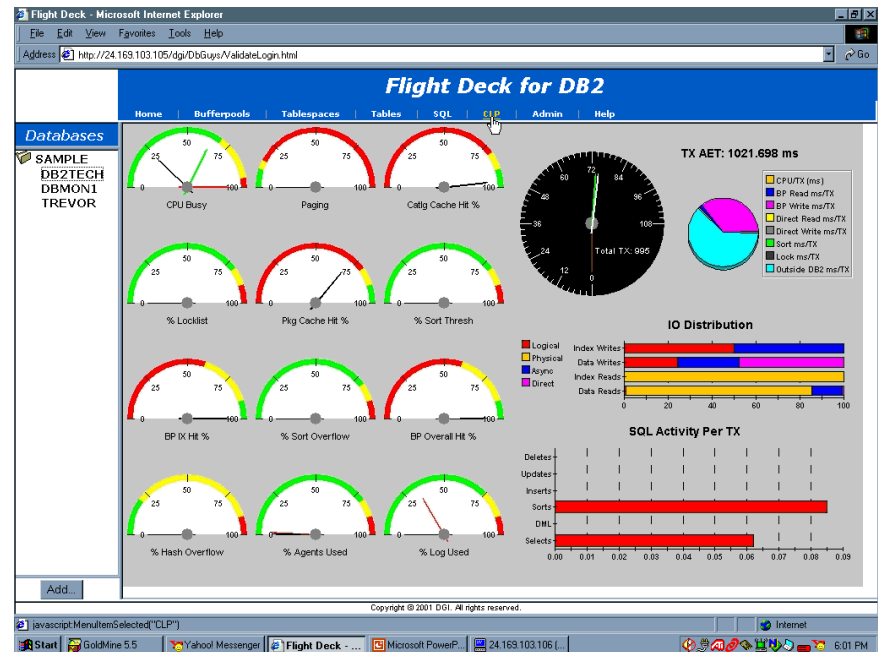


DB2 for Linux, UNIX, Windows

ADJUSTING THE KNOBS: DBM AND DB CFG TUNING FOR DB2 UDB FOR Linux, UNIX and Windows

Philip K. Gunning
Performance Specialist
DGI



- DB2 TM 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.



Agenda

- Introduction
- Why Tune?
- Memory and how it is used.
- Configuration Parameters
- Agent Related Parameters
- Conclusion



Introduction

- Successful tuning requires knowledge of DB2 processing and available monitoring facilities



DB2 UDB Process Model

Client




Applications


Agents



Memory



Prefetchers







Logger



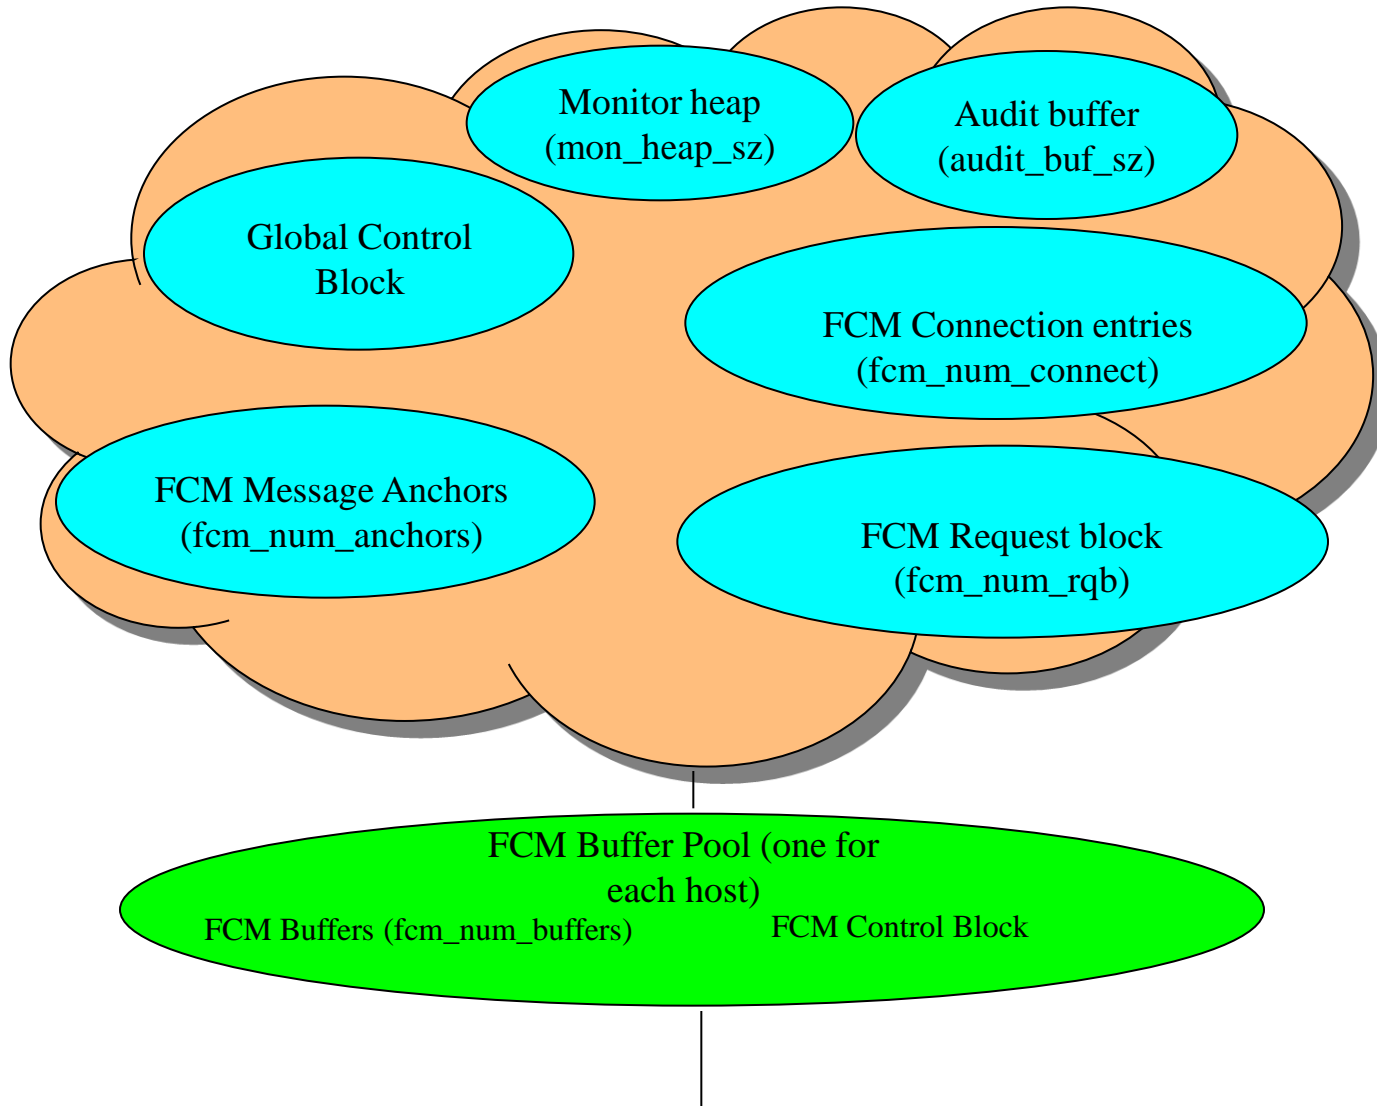
IO Cleaners



Shared Sorts



Database Manager Shared Memory



Database Global Memory



Database Manager Shared Memory

Database Global Memory

Utility Heap (util_heap_sz)	Buffer Pools (buffpage)	Database Heap (dbheap)
Backup Buffer (backbufsz)	Extended Memory Cache	Log Buffer (logbufsz)
Restore Buffer (restbufsz)	Lock List (locklist)	Catalog Cache (catalogcache_sz)
Package Cache (pckcachesz)	Sort Heap – Shared Sort (sortheap)	

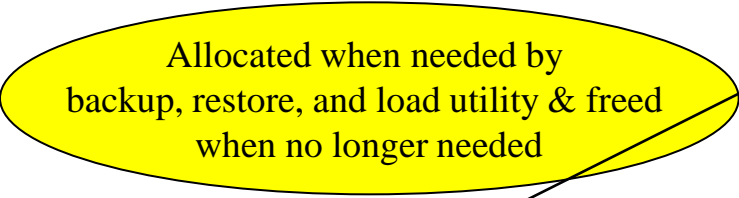






Low

Database Manager Shared Memory

Database Global Memory



Database Manager Shared Memory

Database Global Memory

1,024 4k pages with range of 8 – 524,288
4k pages



Database Global Memory

Backup Buffer
(backbufsz)

Allocated when backup utility is called & freed when the utility completes processing.



Database Manager Shared Memory

Database Global Memory

Allocated when the utility is called and freed when utility completes

Med

Restore Buffer
(restbufsz)

Database Global Memory

1,024 4k pages with range of 16 – 524,288
4k pages



Database Manager Shared Memory



High

Database Global Memory

8x maxappls or 32, whichever is largest with upper limit of 64,000 (32 bit) or 524,288 (64 bit) 4k pages depending on OS

Package Cache
(pckcachesz)

Database Global Memory

Allocated when the database is initialized and when the database is shutdown

Must be large enough to hold all SQL statements that are executing concurrently. Package cache reduces overhead by eliminating the need to access catalog and by eliminating a prepare or the load of a package

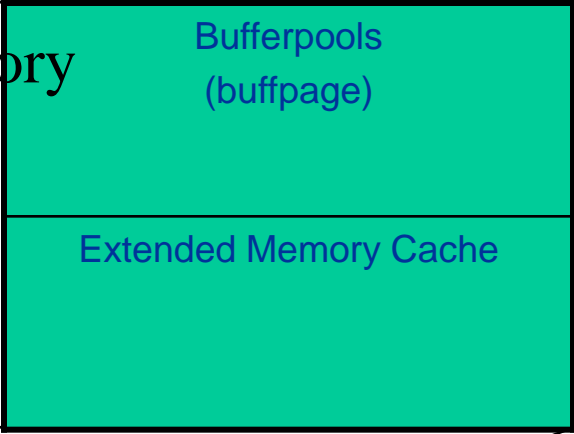




Server Shared Memory

Good use of bufferpools can give you the biggest bang for the buck. Can offset bad design to some extent

Database Global Memory



50 - 75% of physical memory can be devoted to bufferpools if dedicated database server

Database Global Memory

Recent presentation indicated one of the top reasons for poor performance is using default bufpage

Use of extended memory cache can be beneficial when the amount of real memory available exceeds the addressability of the OS, workload is mostly read only, and when the workload is IO bound.





Database Manager Shared Memory

Database Global Memory

Allocated at first connect and freed when last application disconnects from database



Default depends on OS with range of 4 – 60,000 4k pages

Used in conjunction with maxlocks. Maxlocks specifies percent of locklist any one application can use before escalation takes place



Database Manager Shared Memory

Database Global Memory

Allocated at first connect and freed when last application disconnects from the database.

Database Global Memory

Database Heap
(dbheap)

Default depends on the OS with a range of 32 – 524,288 4k pages

Log Buffer, bufferpool control blocks, and Catalog Cache are allocated from dbheap

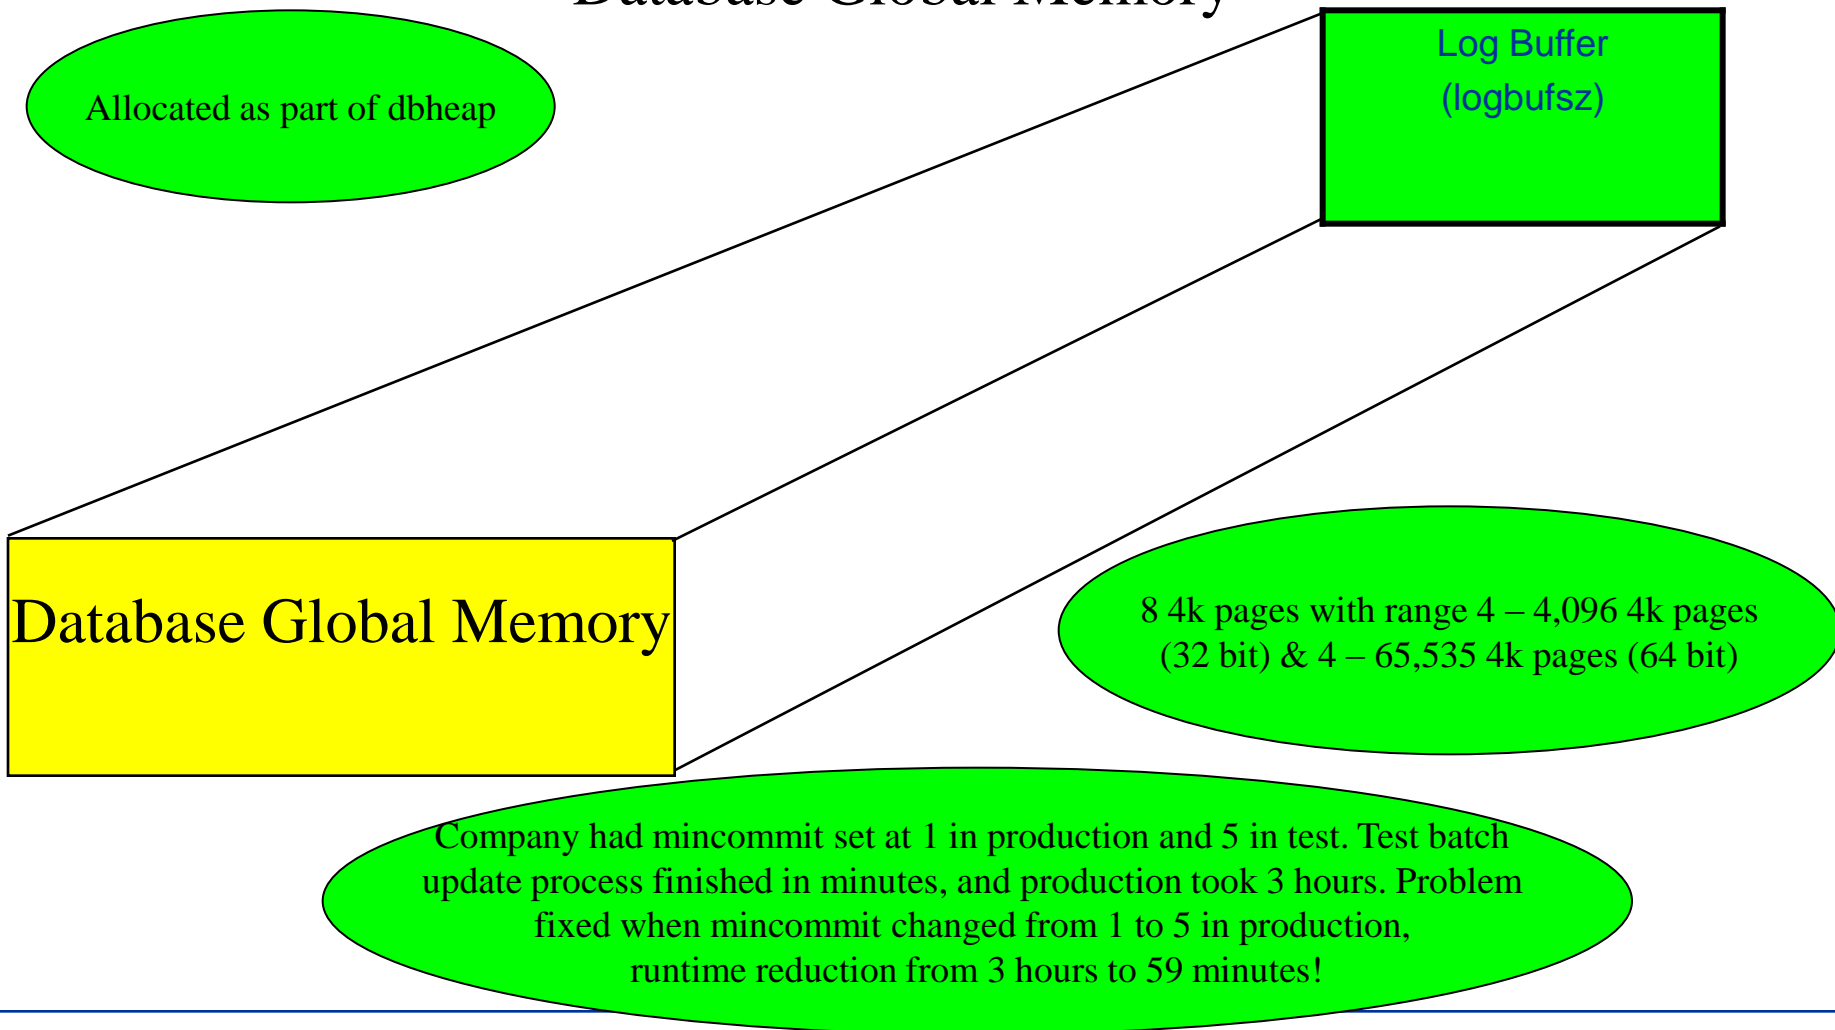




High

Database Manager Shared Memory

Database Global Memory





Med

Database Manager Shared Memory

Database Global Memory

Default depends on OS with range of
16 – 60,000 4k pages

Catalog Cache
(catalogcache_sz)

Database Global Memory

If not large enough increase a few
pages at a time

Stores table descriptor info used when
table, view, or alias referenced
during compilation of an SQL
statement

Monitor using elements:
cat_cache_lookups, cat_cache_inserts,
cat_cache_overflows,
cat_cache_heap_full



Database Manager Shared Memory



Database Global Memory

Sort Heap – Shared Sort
(sortheap)

Database Global Memory

SHEAPTHRES is an instance wide soft limit
for private sorts

SHEAPTHRES for shared sorts is an instance
wide hard limit on the on total amount
of memory used by shared sorts at any
given time



Application Global Memory

(app_ctl_heap_sz)

Med

Only allocated if
if you are using
EEE or EE with intra_parallel
enabled

Agent Private Memory

Used to store Declared Temporary
Tables in EEE

**Application
Heap
(applheapsz)**

**Agent Stack
(agent_stack_sz)**

**Statistics Heap
(stat_heap_sz)**

**Sort Heap
(sortheap)**

**DRDA Heap
(drda_heap_sz)**

**UDF Memory
(udf_mem_sz)**

**Statement Heap
(stmtheap)**

Query Heap (query_heap_sz)

Client I/O Block (rqrioblk)





Agent Private Memory

Application

Heap

(applheapsz)

Allocated when agent initialized and freed when agent completes work for an application. Stores copies of executing SQL statements

Default of 128 or 64 4k pages depending on EEE or not with a range of 16 – 60,000 4k pages



Agent Private Memory

Application

Heap

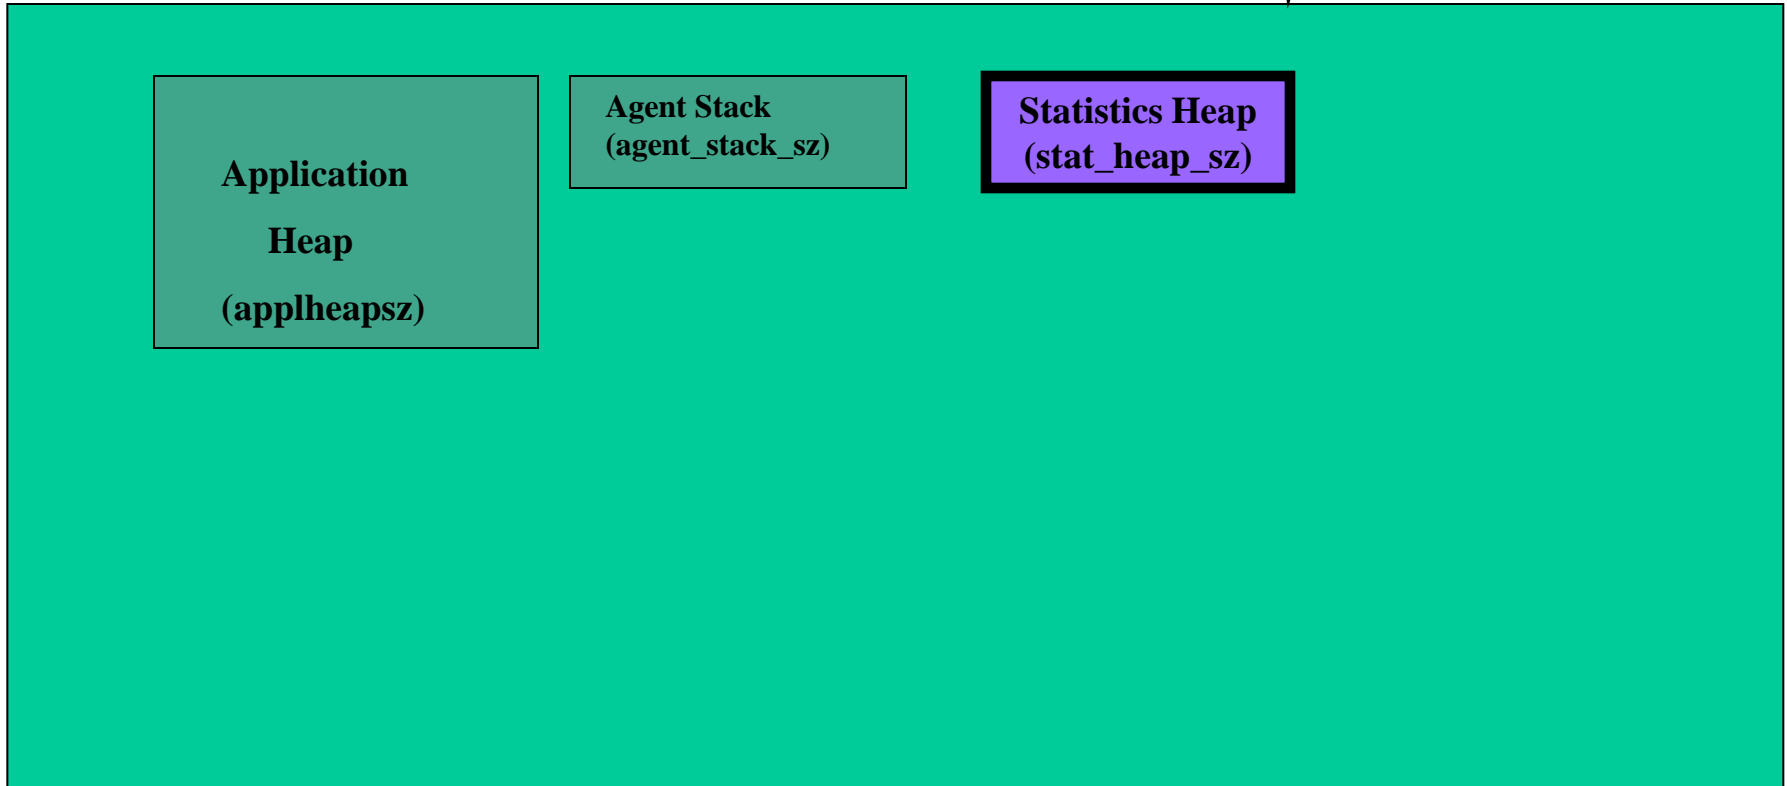

(applheapsz)

Agent Stack
(agent_stack_sz)

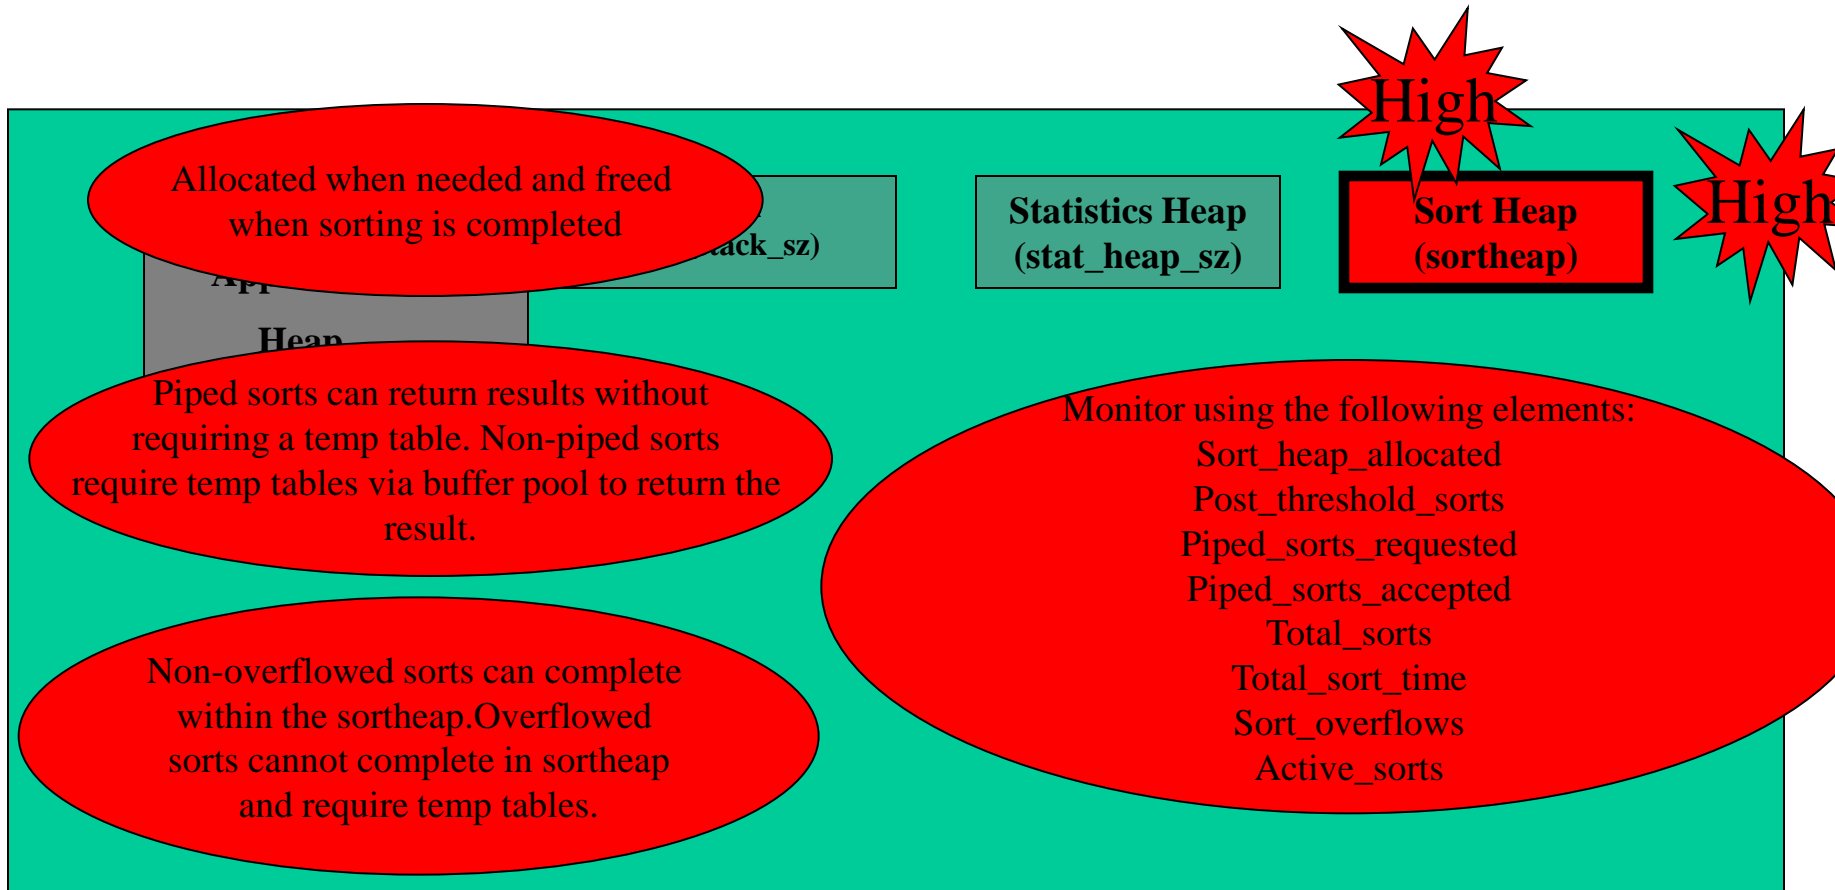
Low unless over-allocated then OS
paging may occur



Agent Private Memory



Agent Private Memory




Measuring SORT Performance

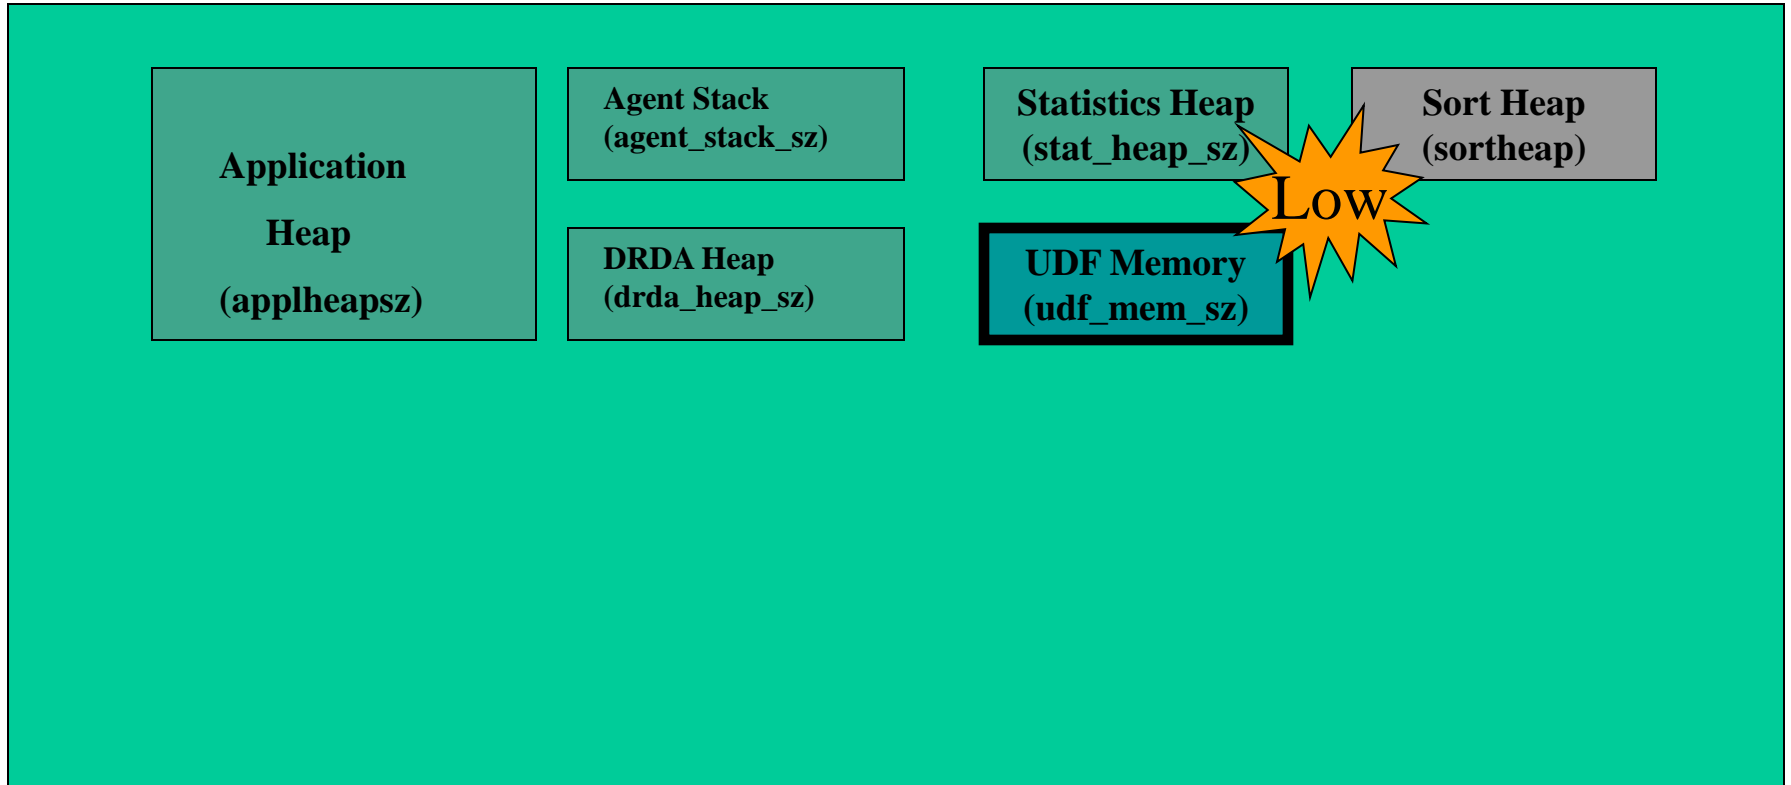
- Connection Events provide **true cost** of Sort activity for an Application
 - create event monitor DGICONN for connections write to file 'e:\tmp\dbaudit\conn\' maxfiles 1 maxfilesize 1024 blocked replace manualstart;
 - Set event monitor DGICONN state = 1;
 - Set event monitor DGICONN state = 0;
 - SQL Events provide **true cost** of Sort activity for an **individual statement**
 - create event monitor DGISQL for statements write to file 'e:\tmp\dbaudit\sql\' maxfiles 1 maxfilesize 2048 blocked replace manualstart;
 - set event monitor DGISQL state = 1;
 - set event monitor DGISQL state = 0;
- db2evmon -path e:\tmp\dbaudit\DIRNAME



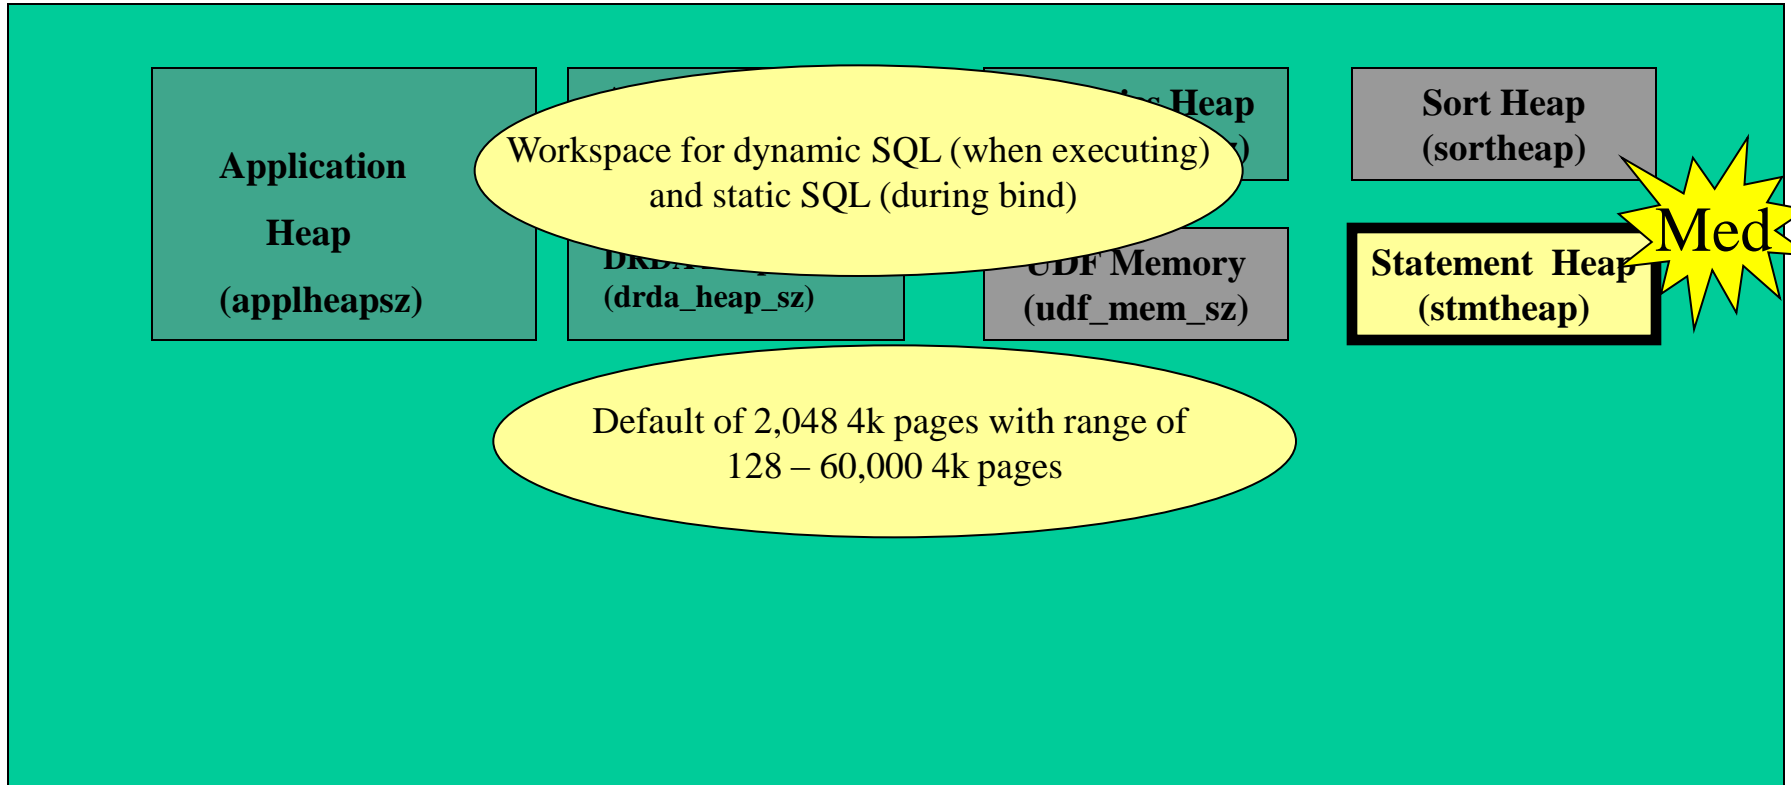
Agent Private Memory




Agent Private Memory



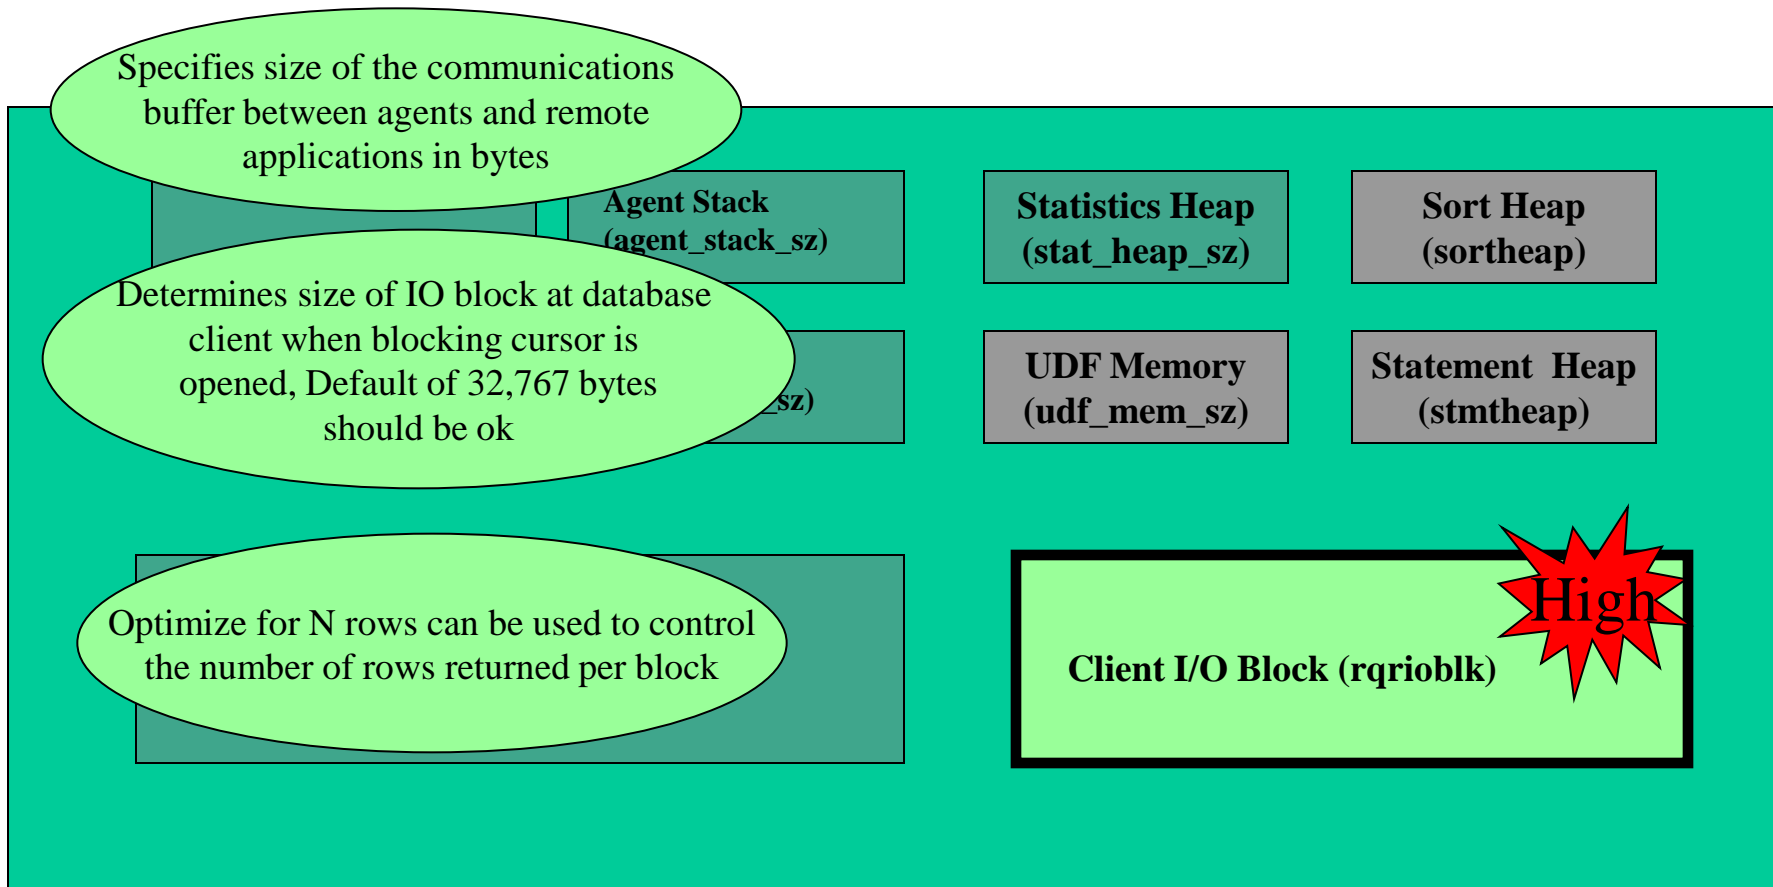
Agent Private Memory



Agent Private Memory



Agent Private Memory



Monitoring

- IBM Self – Managing and Resource Tuning (SMART) Project
- Will make DBAs more productive and job easier
 - DBA will always be in-the-loop
- Compete better with SQL Server and to a lesser extent Oracle



Configure Performance Wizard

Review the performance configuration recommendations.

Based on your selections in this wizard, as well as the volume of data in the database, and system information, this wizard recommends the following values. Below the list, specify if you want to save the new values to a script, or apply them to the database immediately.

Parameter	Current value	Suggested value	DB2 Parameter
Application	128	128	app_cf_heap_sz
Buffer pool ...	250	3456	bufpage
Catalog cac...	32	179	catalogcache_sz
Changed p...	60	60	chngpgs_thresh
Database h...	600	811	dbheap
Default deg...	1	1	df_degrec
Default pref...	16	32	df_prefetch_sz
Maximum s...	50	334	locklist
Log buffer s...	8	32	logbufsz
Log file size	900	250	logflsiz
Number of ...	3	6	logprimary
Number of ...	2	10	logsecond
Maximum n...	40	40	maxappls
Maximum L...	22	18	maxlocks

Apply these recommendations immediately

Save these recommendations to the Script Center

Script name:

Script description: Configure Performance wizard recommendations

◀ Back Finish Cancel

Monitoring

- In order to make sure that resources are being used efficiently and to ensure that business requirements are met, continuous monitoring must be practiced
- Use Snapshot Monitoring and Event Monitoring
- Third Party Vendor Tools



Monitoring

Microsoft PowerPoint - [PittsburghDB2UserGroup]

File Edit View Insert Format Tools Slide Show Window Help

Show Monitor - localhost.DB2.BASEBALL

Monitor Selected View Tools Help

Monitor name: phillocking Status: Started Regular sample: Every 10 Seconds

Description:

Performance variable	Description	Level	Category	Value	Average
Deadlocks	Deadlocks detected	Datab...	Lock and ...	0.0	
Lock Waits	Lock waits during the interval	Datab...	Lock and ...	0.0	
Lock Timeouts	Lock timeouts during the inter...	Datab...	Lock and ...	0.0	

Details Graph Thresholds Graph Settings

Thresholds

What: Applications in Lock wait, Total Exclusive Lock Escalatio, Lock List In Use (bytes)

Last value, Average, Maximum, Minimum, Time

Sep 6, 2001 12:01:10 AM 12:01:20 AM 12:01:30 AM 12:01:40 AM

Sep 6, 2001 12:00:59 AM Sep 6, 2001 12:01:51 AM

Monitor data is refreshed. (12:01:51 AM)

Slide 30 of 42 standardquest

Start Microsoft PowerPoint - [Pi...] Show Monitor - localh... DB2 CLP 12:01 AM



ENVIRONMENTAL VARIABLES

- db2set DB2_BINSORT=YES (default)
- “Enables a new sort algorithm that reduces the CPU time and elapsed time of sorts. This new algorithm extends the extremely efficient integer sorting technique of DB2 UDB to all sort data types such as BIGINT, CHAR, VARCHAR, FLOAT, and DECIMAL, as well as combinations of these data types.”
Administration Guide, Appendix D - Registry Variables, Table 69, Performance.
- **AIX Only**



RAID-5 (6 + 1) EXAMPLE

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=*



RAID-5 (6 + 1) EXAMPLE

db2set DB2_STRIPED_CONTAINERS=


NOT ENABLED IS DEFAULT!

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!



NUM_IOCLEANERS

- DB CFG Default 1 Range(0 –255)
 - Specifies number of asynchronous page cleaners for a database.
 - Write changed pages from Bufferpool to disk
 - Triggered by CHNGPGS_THRESH which specifies a percentage of used pages at which asynchronous page cleaners will start writing out pages
 - Set this to the number of physical devices



NUM_IOSERVERS

- DB CFG Default 3 Range(1 –255)
 - Used to specify the number of prefetchers that work on behalf of database agents to perform prefetch IO and asynchronous IO for utilities such as backup and restore.
 - Set to 1 less than the number of physical disks available



Agent Parameters

- 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



Agent Parameters

- 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



Agent Parameters

- 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



Agent Parameters

- 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



Agent Parameters

- 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



Agent Parameters

- 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



Conclusion

- Successful system tuning requires knowledge of:
 - DB2 processing
 - Available monitoring facilities
 - Instance Configuration Parameters
 - Database Configuration Parameters
 - Cause and Effect of parameters to processing
- *Available References



References

- DB2 UDB V7.1 Performance Tuning Guide
SG24-6012-00
- DB2 UDB V7.1 Administration Guide : Performance,
SC09-2945-00
- DB2 UDB System Monitor Guide and Reference,
SC09-2956-00
- DB2 UDB Administration Guide: Planning
SC09-2944-00



References

- DB2 UDB V7.1 Command Reference, SC09-2951-00
- Database Performance on AIX in DB2 UDB and Oracle Environments, SG24-5511-00
- URL: <http://www.software.ibm.com/data/db2/udb>



Thank you!

Phil Gunning

DGI

pgunning@

breakthroughdb2

.com

877-4DB-GUYS X 8

