



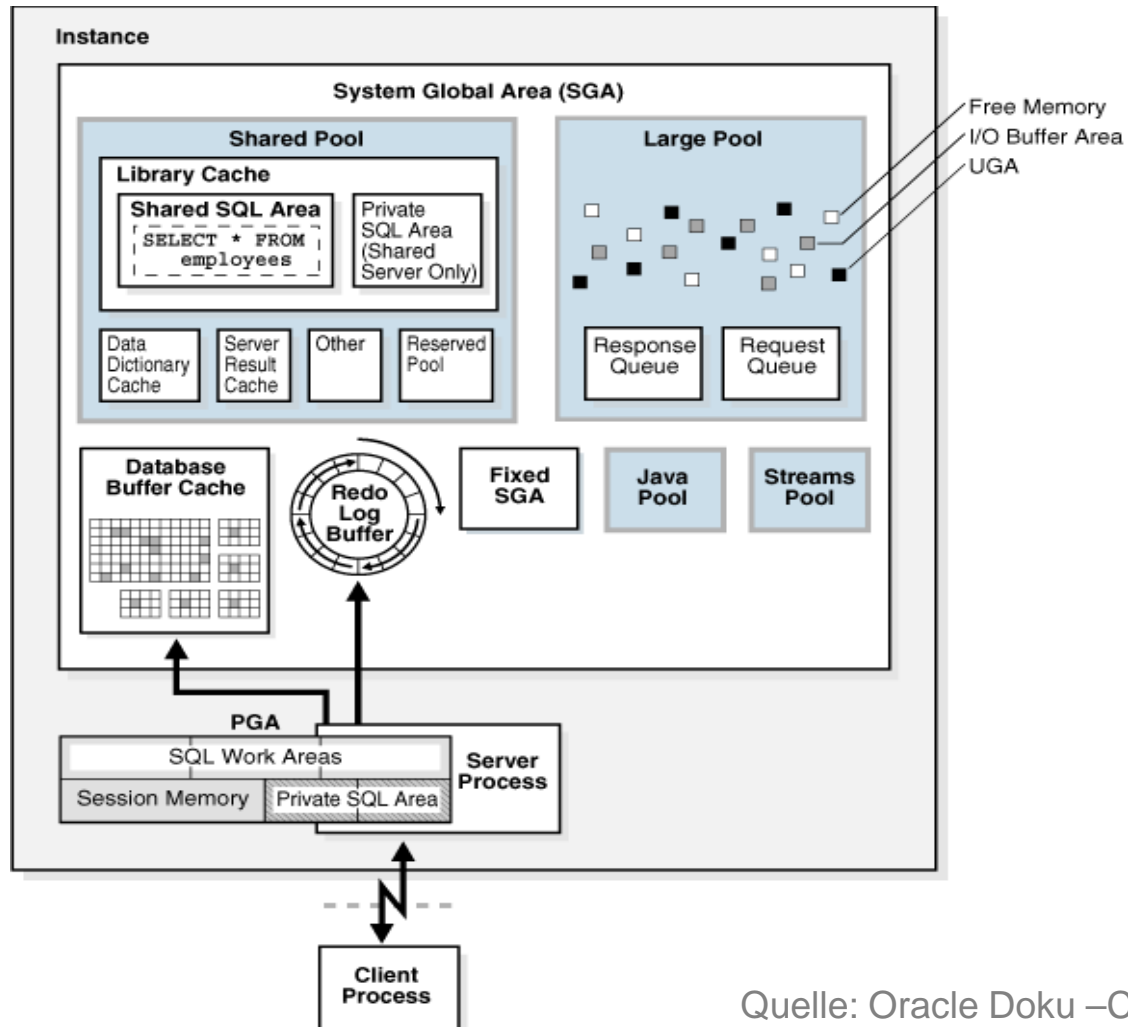
Memory-Drilldown von der SGA über die PGA zum Database Buffer Advisor

DOAG Konferenz
20. - 22.11.2012

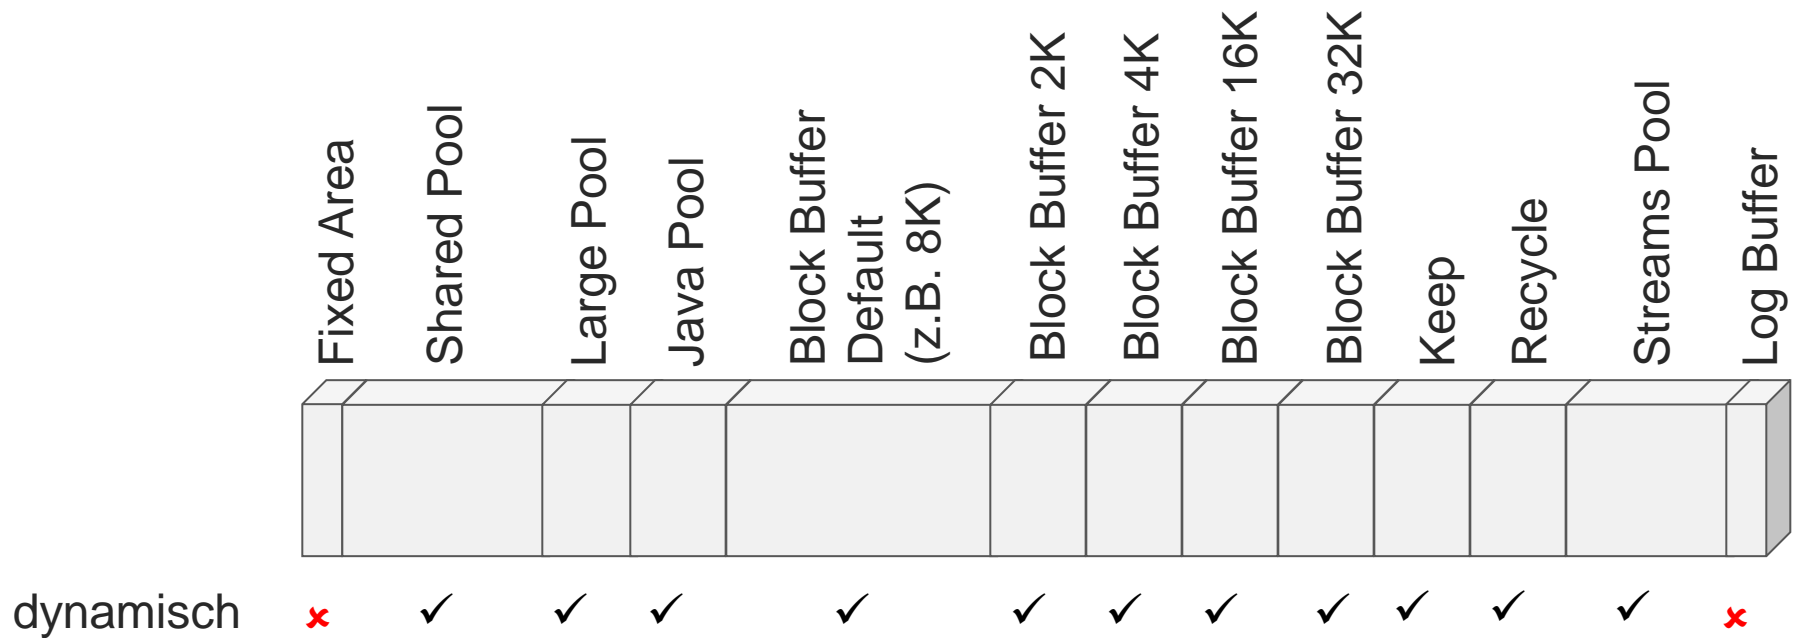
Klaus Reimers
kr@ordix.de
www.ordix.de

- SGA
 - Variable Size
 - Shared Pool
 - Large Pool
 - Java Pool
 - Streams Pool
 - Database Buffer
 - Log Buffer
 - Dynamische SGA
- PGA
- Dynamische Memory-Verwaltung
- Advisor
 - PGA Advisor
 - Database Buffer Advisor

Überblick SGA (I)



Quelle: Oracle Doku –Concept Guide



SGA - Infos

```
SQL> show sga;
Total System Global Area  289406976 bytes
Fixed Size                 1290184 bytes
Variable Size             121634872 bytes
Database Buffers         159383552 bytes
Redo Buffers              7098368 bytes
```

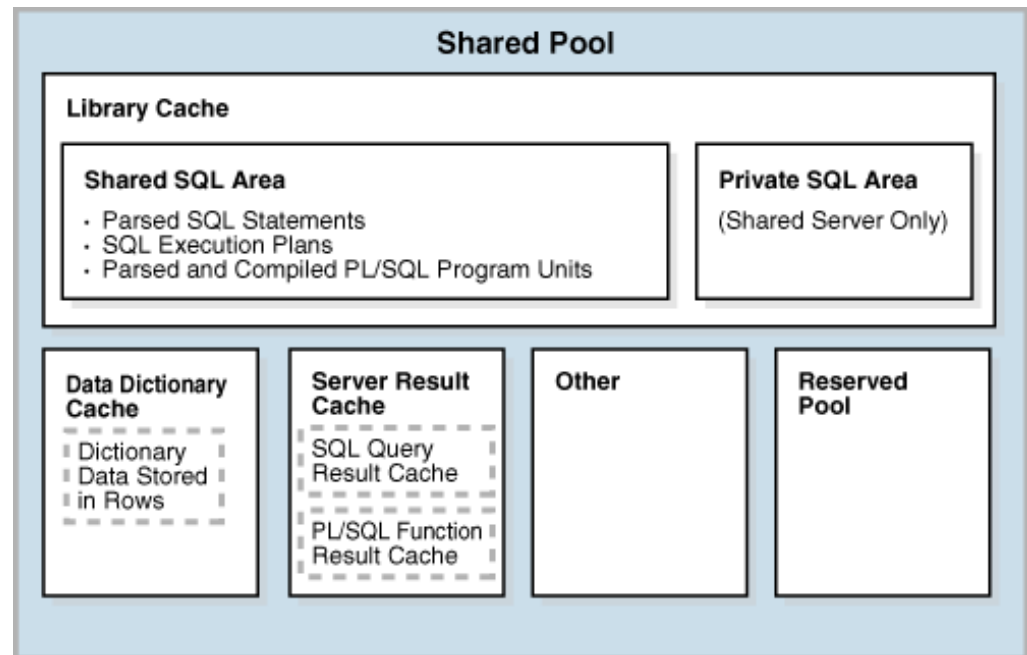
```
SQL> select * from v$sgainfo;
NAME                                BYTES RES
-----
Fixed SGA Size                      1290184 No
Redo Buffers                        7098368 No
Buffer Cache Size                   159383552 Yes
Shared Pool Size                    79691776 Yes
Large Pool Size                     33554432 Yes
Java Pool Size                      4194304 Yes
Streams Pool Size                   4194304 Yes
Granule Size                        4194304 No
Maximum SGA Size                    289406976 No
Startup overhead in Shared Pool     37748736 No
Free SGA Memory Available           0
```

- SGA
 - Variable Size
 - Shared Pool
 - Large Pool
 - Java Pool
 - Streams Pool
 - Database Buffer
 - Log Buffer
 - Dynamische SGA
- PGA
- Dynamische Memory Verwaltung
- Advisor
 - PGA Advisor
 - Database Buffer Advisor

Shared Pool

Überblick

- Initialisierungsparameter :
 - shared_pool_size
- Views
 - v\$librarycache
 - v\$rowcache



Quelle: Oracle Doku – Concept Guide

Instance Efficiency Percentages (Target 100%)

```

~~~~~
          Buffer Nowait %:    92.20          Redo NoWait %:    99.93
          Buffer Hit %:      99.60          In-memory Sort %: 100.00
          Library Hit %:    97.77          Soft Parse %:    88.31
          Execute to Parse %: 79.60          Latch Hit %:     97.27
Parse CPU to Parse Elapsed %: 63.01          % Non-Parse CPU:  97.54
    
```

Namespace	Get Requests	Pct Miss	Pin Requests	Pct Miss	Reloads	Invali-dations
BODY	12,062	0.0	12,062	0.0	0	0
CLUSTER	93	0.0	81	0.0	0	0
INDEX	1,914	0.2	1,072	0.3	0	0
SQL AREA	131,290	17.7	1,468,740	3.3	977	2
TABLE/PROCEDURE	79,815	0.0	708,512	0.1	281	0
TRIGGER	3,138	0.1	3,138	0.2	3	0

Dictionary Cache

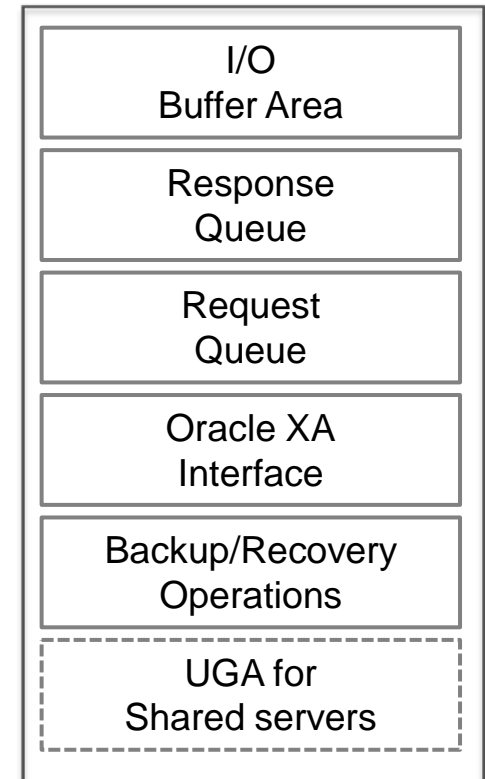
Cache	Get Requests	Pct Miss	Scan Reqs	Pct Miss	Mod Reqs	Final Usage
dc_files	120	0.0	0		0	30
dc_histogram_data	86,799	0.0	0		0	3,363
dc_histogram_data_values	1,350	0.0	0		0	537
dc_histogram_defs	1,247,234	0.0	0		0	4,083
dc_object_ids	337,795	0.0	0		0	792
dc_objects	23,879	0.1	0		0	1,228
dc_profiles	1,168	0.0	0		0	1
dc_rollback_segments	15,848,525	0.0	0		0	957
dc_segments	114,716	0.4	0		33	3,476
dc_sequences	23,514	75.8	0		23,515	22
dc_tablespace_quotas	33	39.4	0		33	0
dc tablespaces	25,330	0.0	0		0	9
dc_user_grants	462	0.0	0		0	24
dc usernames	540	0.0	0		0	9
dc_users	36,012	0.0	0		0	28

- SGA
 - Variable Size
 - Shared Pool
 - Large Pool
 - Java Pool
 - Streams Pool
 - Database Buffer
 - Log Buffer
 - Dynamische SGA
- PGA
- Dynamische Memory-Verwaltung
- Advisor
 - PGA Advisor
 - Database Buffer Advisor

- wird benötigt bei:
 - paralleler Verarbeitung
 - Multi Threaded Server
 - RMAN über mehrere Kanäle
- Parameter
 - large_pool_size

```
SQL> select * from v$sgastat
      where pool like '%large%';
```

POOL	NAME	BYTES
large pool	PX msg pool	3894304
large pool	free memory	300000



Quelle:
oraclegurus.wordpress.com

- SGA
 - Variable Size
 - Shared Pool
 - Large Pool
 - Java Pool
 - Streams Pool
 - Database Buffer
 - Log Buffer
 - Dynamische SGA
- PGA
- Dynamische Memory-Verwaltung
- Advisor
 - PGA Advisor
 - Database Buffer Advisor

- wird benötigt bei:
 - Verwendung von Java in der Datenbank
 - anstelle von PL/SQL
- Parameter
 - java_pool_size
- Views
 - V\$JAVA_LIBRARY_CACHE_MEMORY
 - V\$JAVA_POOL_ADVICE

```
SQL> select * from v$sgastat
        where pool like '%java%';
```

POOL	NAME	BYTES
-----	-----	-----
java pool	free memory	4194304

- SGA
 - Variable Size
 - Shared Pool
 - Large Pool
 - Java Pool
 - Streams Pool
 - Database Buffer
 - Log Buffer
 - Dynamische SGA
- PGA
- Dynamische Memory-Verwaltung
- Advisor
 - PGA Advisor
 - Database Buffer Advisor

- wird benötigt bei:
 - Verwendung von Advanced Queues (AQ)
 - Streams Replication
 - Logical Data Guard
 - Audit Vault
- Parameter
 - streams_pool_size
- Views
 - V\$STREAMS_POOL_ADVICE
 - diverse andere Views
 - v\$streams_....

```
SQL> select * from v$sgastat
        where pool like '%streams%';
```


POOL	NAME	BYTES
streams	pool free memory	4194304

- SGA
 - Variable Size
 - Shared Pool
 - Large Pool
 - Java Pool
 - Streams Pool
 - Database Buffer
 - Log Buffer
 - Dynamische SGA
- PGA
- Dynamische Memory-Verwaltung
- Advisor
 - PGA Advisor
 - Database Buffer Advisor

- Ein oder mehrere Puffer zum Caching von DB-Blöcken
- Verwaltung der Puffer durch einzelne LRU-Listen
- Lesen der Blöcke über Benutzerprozesse
- Schreiben der Blöcke durch Database Writer



Beispiel: Default-Blockgröße 8K, Verwendung aller Buffer Pools



- v\$dbuffer_pool

```
SQL> select id, name, block_size, current_size from v$dbuffer_pool;
```

ID	NAME	BLOCK_SIZE	CURRENT_SIZE
1	KEEP	8192	12
2	RECYCLE	8192	12
3	DEFAULT	8192	128
4	DEFAULT	2048	12
5	DEFAULT	4096	12
7	DEFAULT	16384	12


- statspack / AWR

P	Number of Buffers	Cache Hit %	Buffer Gets	Physical Reads	Physical Writes	Free Buffer Waits	Write Complete Waits	Buffer Busy Waits
D	371,122	99.6	64,297,973	262,145	1,233,928	0	14	284,849
K	24,869	100.0	576,376	157	2,110	0	0	64

- Verhältnis von logischem zu physikalischem I/O
- statspack / AWR

```
Instance Efficiency Percentages (Target 100%)
~~~~~
      Buffer Nowait %:    92.20          Redo NoWait %:    99.93
      Buffer Hit %:     99.60          In-memory Sort %: 100.00
      Library Hit %:   97.77          Soft Parse %:    88.31
      Execute to Parse %: 79.60          Latch Hit %:    97.27
      Parse CPU to Parse Elapsd %: 63.01          % Non-Parse CPU: 97.54
```

- SGA
 - Variable Size
 - Shared Pool
 - Large Pool
 - Java Pool
 - Streams Pool
 - Database Buffer
 - Log Buffer
 - Dynamische SGA
- PGA
- Dynamische Memory-Verwaltung
- Advisor
 - PGA Advisor
 - Database Buffer Advisor



Logging (Events)

Event	Waits	Time (s)	Avg wait (ms)	%Total Call Time
PL/SQL lock timer	739	3,492	4726	45.5
CPU time		2,633		34.3
log file parallel write	82,362	428	5	5.6
log file sync	34,884	327	9	4.3
db file sequential read	124,127	286	2	3.7

Event	Total Waits	% of Waits							
		<1ms	<2ms	<4ms	<8ms	<16ms	<32ms	<=1s	>1s
Log archive I/O	70	12.9	4.3	4.3	1.4		1.4	74.3	1.4
log buffer space	15						33.3	66.7	
log file parallel write	82K	9.6	30.8	24.8	20.1	10.6	3.0	1.1	
log file sequential read	1174	20.1	.4	8.3	30.2	20.4	6.5	14.1	.1
log file single write	64	67.2	26.6	4.7	1.6				
log file switch (private s	26				11.5	46.2	30.8	11.5	
log file sync	34K	8.1	19.4	25.5	25.0	13.1	5.8	3.3	...


- SGA
 - Variable Size
 - Shared Pool
 - Large Pool
 - Java Pool
 - Streams Pool
 - Database Buffer
 - Log Buffer
 - Dynamische SGA
- PGA
- Dynamische Memory-Verwaltung
- Advisor
 - PGA Advisor
 - Database Buffer Advisor

	Fixed Area	Shared Pool	Large Pool	Java Pool	Block Buffer Default (z.B. 8K)	Block Buffer 2K	Block Buffer 4K	Block Buffer 16K	Block Buffer 32K	Keep	Recycle	Streams Pool	Log Buffer
dynamisch	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	x
Anpassung durch ASMM	-	✓	✓	✓	✓	x	x	x	x	x	x	✓	-

ASMM automatisch durch Setzen von sga_target

- SGA
 - Variable Size
 - Shared Pool
 - Large Pool
 - Java Pool
 - Streams Pool
 - Database Buffer
 - Log Buffer
 - Dynamische SGA
- PGA
- Dynamische Memory-Verwaltung
- Advisor
 - PGA Advisor
 - Database Buffer Advisor

- Größe zum Betriebssystem dynamisch, intern dynamisch
- `PGA_AGGREGATE_TARGET = integer [K | M | G]`
- Statistiken in `v$sysstat`
 - `WORKAREA MEMORY ALLOCATED`
 - `WORKAREA EXECUTIONS - OPTIMAL`
 - `WORKAREA EXECUTIONS - ONEPASS`
 - `WORKAREA EXECUTIONS - MULTIPASS`
- View `v$pgastat`
- View `v$process`
 - `PGA_USED_MEM`
 - `PGA_ALLOCATED_MEM`
 - `PGA_MAX_MEM`

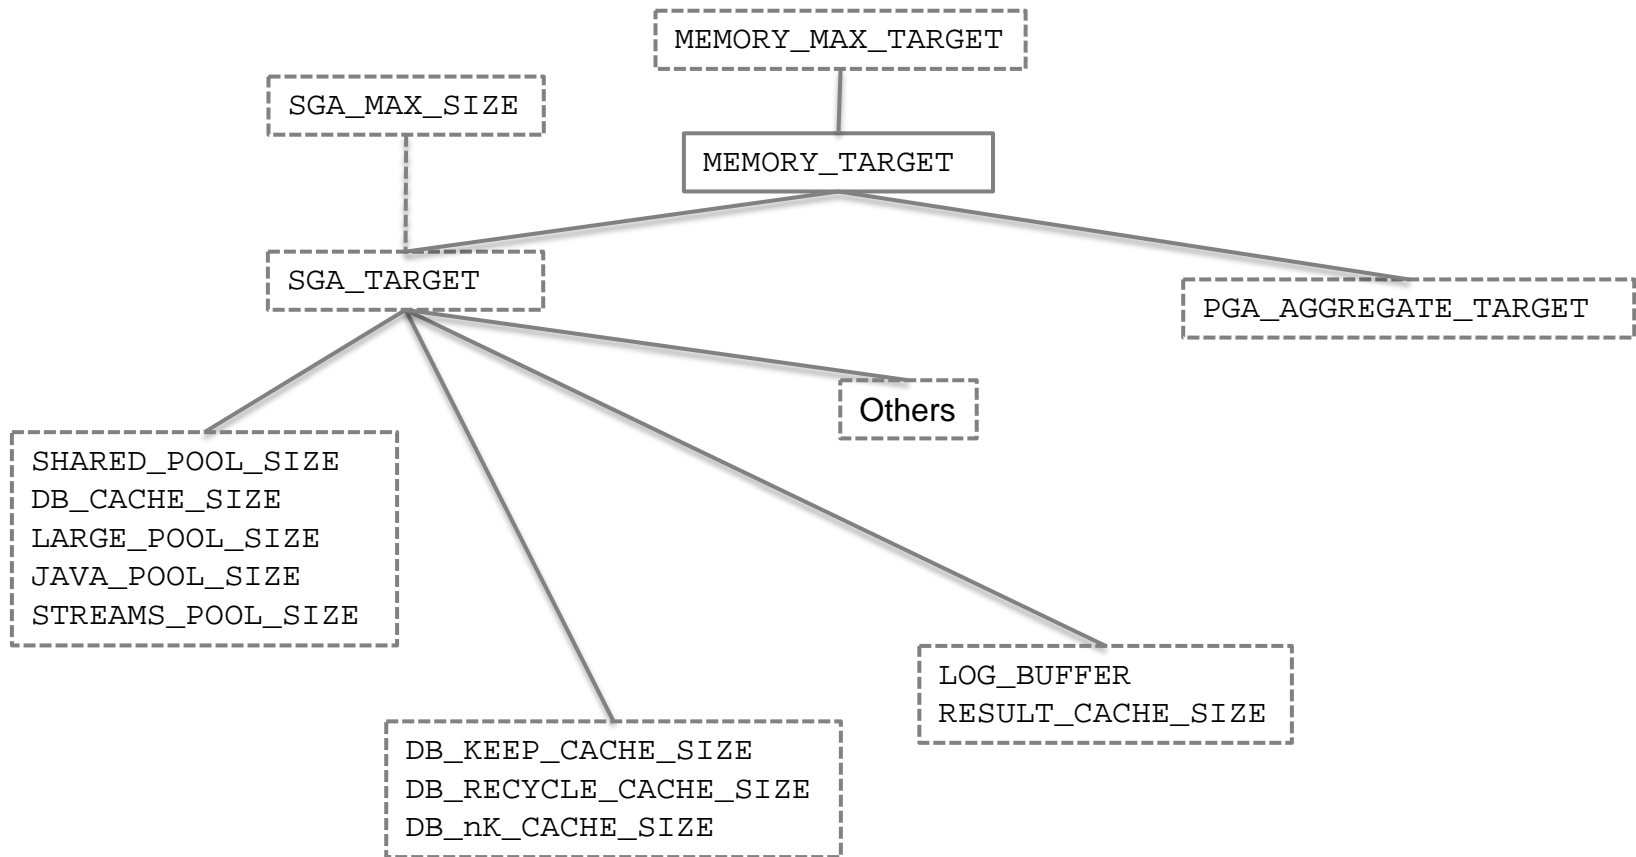


Quelle: Oracle Doku – Concept Guide

- Informationsquelle
 - statspack
 - AWR
- möglichst realistische Größe

	PGA Aggr Target (M)	Auto PGA Target (M)	PGA Mem Alloc (M)	W/A PGA Used (M)	%PGA W/A Mem	%Auto W/A Mem	%Man W/A Mem	Global Mem Bound (K)
B	655	41	2,663.1	4.6	.2	100.0	.0	102,400
E	655	41	2,682.1	5.6	.2	100.0	.0	102,400

- SGA
 - Variable Size
 - Shared Pool
 - Large Pool
 - Java Pool
 - Streams Pool
 - Database Buffer
 - Log Buffer
 - Dynamische SGA
- PGA
- Dynamische Memory-Verwaltung
- Advisor
 - PGA Advisor
 - Database Buffer Advisor



- SGA
 - Variable Size
 - Shared Pool
 - Large Pool
 - Java Pool
 - Streams Pool
 - Database Buffer
 - Log Buffer
 - Dynamische SGA
- PGA
- Dynamische Memory-Verwaltung
- Advisor
 - PGA Advisor
 - Database Buffer Advisor

PGA Aggr Target Est (MB)	Size Factr	W/A MB Processed	Estd Extra W/A MB Read/Written to Disk	Estd Time to Process Bytes (s)	Estd PGA Cache Hit %	Estd PGA Overalloc Count
82	0.1	41,255	21,928	14.5	65.0	1,432
164	0.3	41,255	21,903	14.5	65.0	1,426
327	0.5	41,255	6,990	11.1	86.0	962
491	0.8	41,255	3,918	10.4	91.0	671
655	1.0	41,255	330	9.5	99.0	606
786	1.2	41,255	146	9.5	100.0	518
916	1.4	41,255	146	9.5	100.0	443
1,047	1.6	41,255	146	9.5	100.0	413
1,178	1.8	41,255	146	9.5	100.0	402
1,309	2.0	41,255	146	9.5	100.0	382
1,964	3.0	41,255	146	9.5	100.0	24
2,618	4.0	41,255	146	9.5	100.0	0
3,928	6.0	41,255	146	9.5	100.0	0
5,237	8.0	41,255	146	9.5	100.0	0

Database Buffer Advisor

P	Size for Estimate (M)	Size (M) Factr	Buffers for Estimate	Est Physical Read Factor	Estimated Physical Reads
D	208	.1	25,805	9.95	208,663,869
D	416	.2	51,610	7.86	164,921,128
D	624	.3	77,415	6.56	137,636,252
D	832	.4	103,220	5.90	123,702,518
D	1,040	.5	129,025	5.51	115,621,486
D	1,248	.6	154,830	5.13	107,524,073
D	1,456	.7	180,635	3.88	81,338,384
D	1,664	.8	206,440	3.34	70,052,343
D	1,872	.9	232,245	2.01	42,080,152
D	2,048	1.0	254,080	1.00	20,979,933
D	2,080	1.0	258,050	0.95	19,903,913
D	2,288	1.1	283,855	0.79	16,569,140
D	2,496	1.2	309,660	0.77	16,080,976
D	2,704	1.3	335,465	0.76	15,954,021
D	2,912	1.4	361,270	0.76	15,855,967
D	3,120	1.5	387,075	0.75	15,744,808
D	3,328	1.6	412,880	0.75	15,670,975
D	3,536	1.7	438,685	0.70	14,753,386
D	3,744	1.8	464,490	0.68	14,277,508
D	3,952	1.9	490,295	0.68	14,255,393
D	4,160	2.0	516,100	0.68	14,212,450

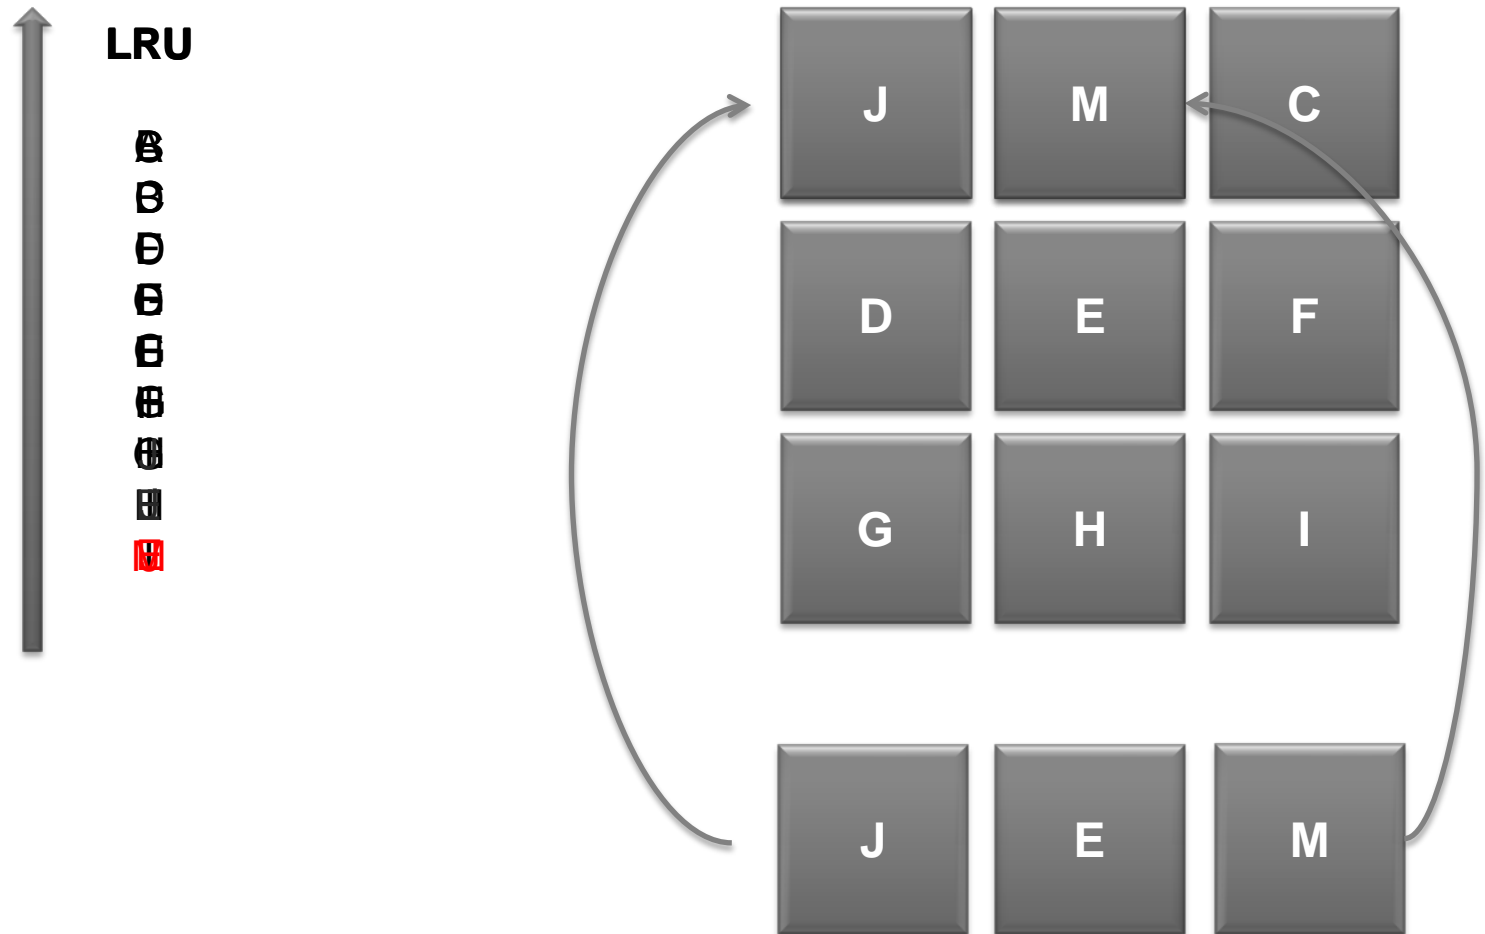
Top 5 Timed Events

~~~~~

| Event                   | Waits      | Time (s) | % Ela Time |
|-------------------------|------------|----------|------------|
| db file scattered read  | 17,443,552 | 42017    | 42.35      |
| db file sequential read | 9,135,007  | 12499    | 10.04      |
| latch free              | 1,156      | 88       | 8.83       |
| buffer busy waits       | 395        | 33       | 3.35       |

# Verdrängung im Database Buffer

Sequential Read



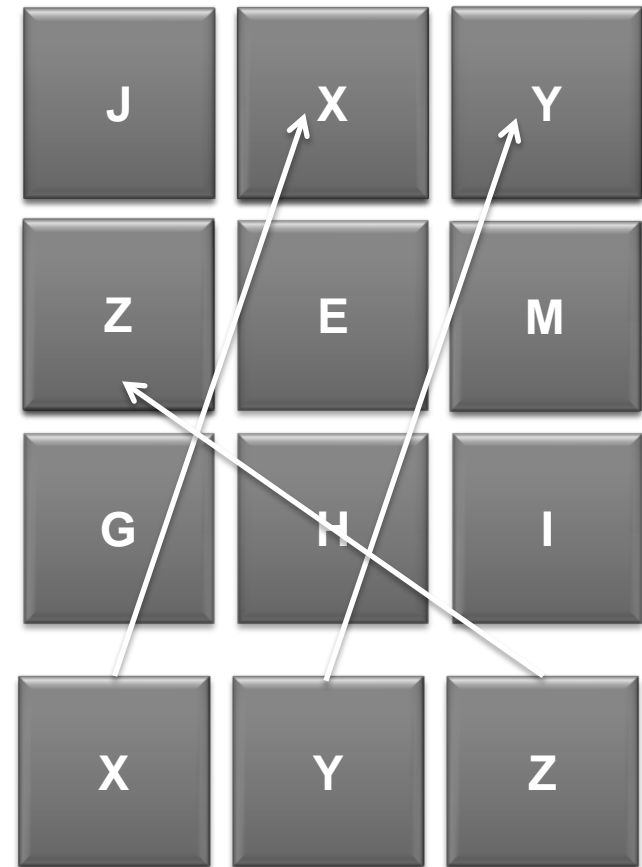

## Top 5 Timed Events

~~~~~

Event	Waits	Time (s)	% Ela Time
db file scattered read	17,443,552	42017	42.35
db file sequential read	9,135,007	12499	10.04
latch free	1,156	88	8.83
buffer busy waits	395	33	3.35

Verdrängung im Database Buffer

Scattered Read



Top 5 Timed Events

~~~~~

| Event                   | Waits      | Time (s) | % Ela Time |
|-------------------------|------------|----------|------------|
| db file scattered read  | 17,443,552 | 42017    | 42.35      |
| db file sequential read | 9,135,007  | 12499    | 10.04      |
| latch free              | 1,156      | 88       | 8.83       |
| buffer busy waits       | 395        | 33       | 3.35       |


```
SQL> alter table mitarbeiter cache;
```

```
SQL> select cache, count(*) from dba_tables group by cache;
```

| CACHE | COUNT(*) |
|-------|----------|
| N     | 8839     |
| Y     | 1        |

# Verdrängung im Database Buffer

Table Cache



- verlässlich
  - fast ausschließliche sequential reads
- nicht verlässlich
  - viele scattered reads (Full Table Scans)
  - wenn Tabellen gecacht werden
  
- alle Advisor hinterfragen
- nicht direkt umsetzen
- Erfolg / Verbesserung verifizieren



Zentrale Paderborn  
Westernmuer 12 - 16  
33098 Paderborn  
Tel.: 05251 1063-0

Seminarzentrum Wiesbaden  
Kreuzberger Ring 13  
65205 Wiesbaden  
Tel.: 0611 77840-00

Zentrales Fax:  
0180 1 67349 0  
0180 1 ORDIX 0

Weitere Geschäftsstellen  
in Köln, Münster und Neu-Ulm

E-Mail: [info@ordix.de](mailto:info@ordix.de)  
Internet: <http://www.ordix.de>

**Vielen Dank für Ihre Aufmerksamkeit!**