

eXstor



TSM Mysteries – Server Memory

Episode 1

Oxford University
TSM Symposium 2007

Michael Abel & Bruno Friess

Agenda



- Why care for memory?
- Types of memory
- TSM memory map
- TSM memory tuning
- Memory layout
- Memory analysis

- For commands, settings, options, values, limites, implications...
- Check documentation valid for OS, TSM server and TSM client level used in your specific environment



Why care for memory?



- TSM Server needs memory
 - ... from some MBs
 - ... up to several GBs
- What we want to know:
 - What is memory used for?
 - Does size really matter?
 - How to tune memory inside TSM
 - How to tune memory on the OS
 - How to get more out of less

What if we ignore...



- ...all those fancy foils?

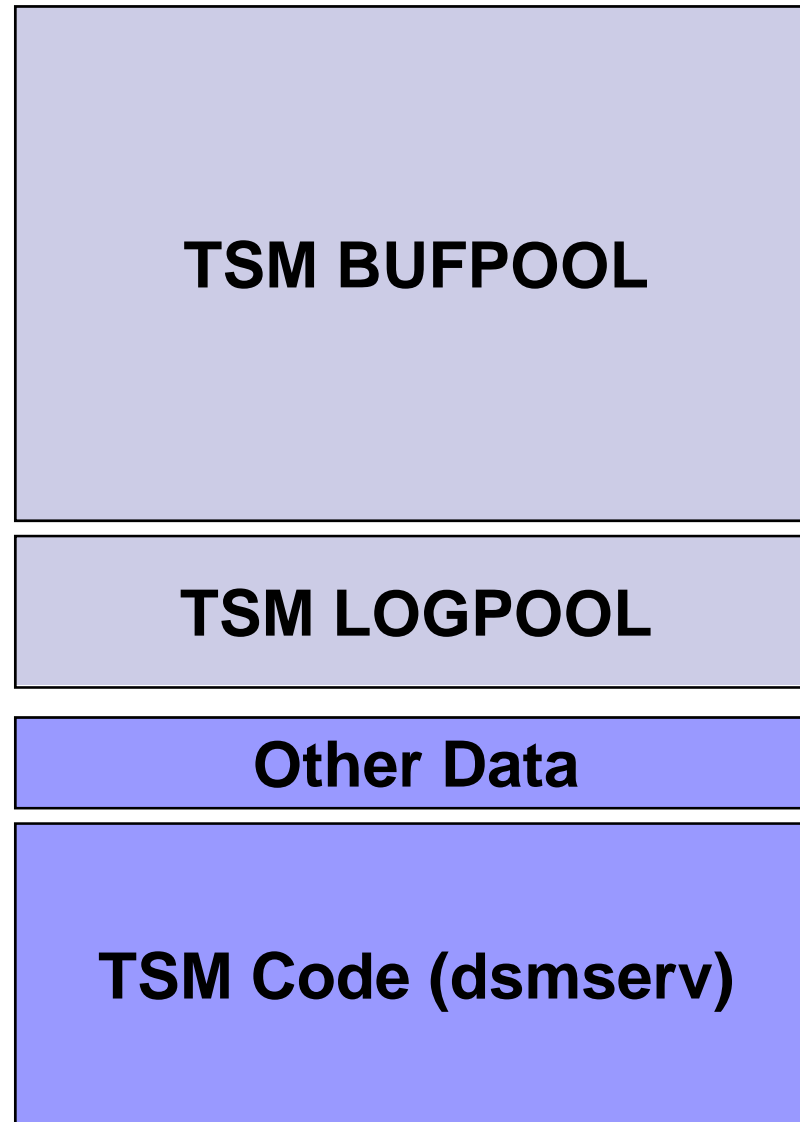
- Worst case scenario:
 - TSM server hogs up all memory
 - OS is running out of memory
 - OS has to page/swap processes
 - Paging creates unwanted IOs
 - Performance of system goes down
 - Performance of TSM goes down
 - TSM admin is in deep trouble

Types of memory



- Most operating systems separate memory into different types
 - Executable binary code
 - Executable (shared) libraries
 - Private data
 - Shared data
 - Kernel memory
 - ...
- Differentiation necessary for
 - Protection, paging, caching, sharing...

TSM memory map



Svmon of dsmserv



Pid	Command	Inuse	Pin	Pgsp	Virtual	64-bit	Mthrd	16MB
282772	dsmserv	19519	5432	0	17975	Y	Y	N

Vsid	Esid	Type	Description	PSize	Inuse	Pin	Pgsp	Virtual
0	0	work	kernel	s	9204	5392	0	9204
5355	11	work	text data BSS heap	s	5945	0	0	5945
1c34c	10	clnt	text data BSS heap, /dev/hd2:61669 dsmserv code	s	1530	0	-	-
1001	9ffffffd	work	shared library	s	1252	0	0	1252
10a0	90000000	work	shared library text	s	988	0	0	988
1e14e	90020014	work	shared library	s	159	0	0	159
8358	9001000a	work	shared library data	s	149	0	0	149
11361	-	work		s	107	33	0	107
4374	80000000	work	private load text	s	51	0	0	51
1f32f	80020014	work	USLA heap	s	33	0	0	33
5375	8001000a	work	private load data	s	22	0	0	22
13343	8fffffff	work	private load data	s	21	0	0	21
4354	ffffffff	work	application stack	s	18	0	0	18
17347	f00000002	work	process private	s	14	7	0	14
14014	9ffffffe	work	shared library	s	12	0	0	12
1a14a	9fffffff	clnt	USLA text, /dev/hd2:2041	s	10	0	-	-
12042	-	clnt	/dev/hd2:14537	s	2	0	-	-
a35a	-	clnt	/dev/lvtsm01:79 admserv.lock	s	1	0	-	-
13363	-	clnt	/dev/lvtsm01:86	s	1	0	-	-
2372	-	clnt	/dev/lvtsm01:4097 diskpool	s	0	0	-	-
14364	-	clnt	/dev/lvtsm01:83 database vol	s	0	0	-	-
15365	-	clnt	/dev/lvtsm01:82 log vol	s	0	0	-	-

- **BUFPOOLSIZE**
 - Set in dsmserv.opt
 - Given in KB
 - Minimum = 256
 - Default = 32768
 - Maximum = virtual memory

- Visible as real memory consumption
- Conservative value 1/8-1/2 of RAM

- LOGPOOLSIZE
 - Set in dsmserv.opt
 - Specified in KB
 - Minimum = 128
 - Default = 512
 - Maximum = virtual memory

- Visible as real memory consumption
- Old max value known 8192
- Performance effect hard to prove

■ AIXASYNCIO

- Server performs DB I/O asynchronous by using AIO interface offered by AIX. Collects smaller IOs into one big IO. Used for databases like Oracle.

■ AIXDIRECTIO

- Bypass AIX file cache when performing IOs to the disk storage pools. Can be specified on the filesystem level also (mount option)

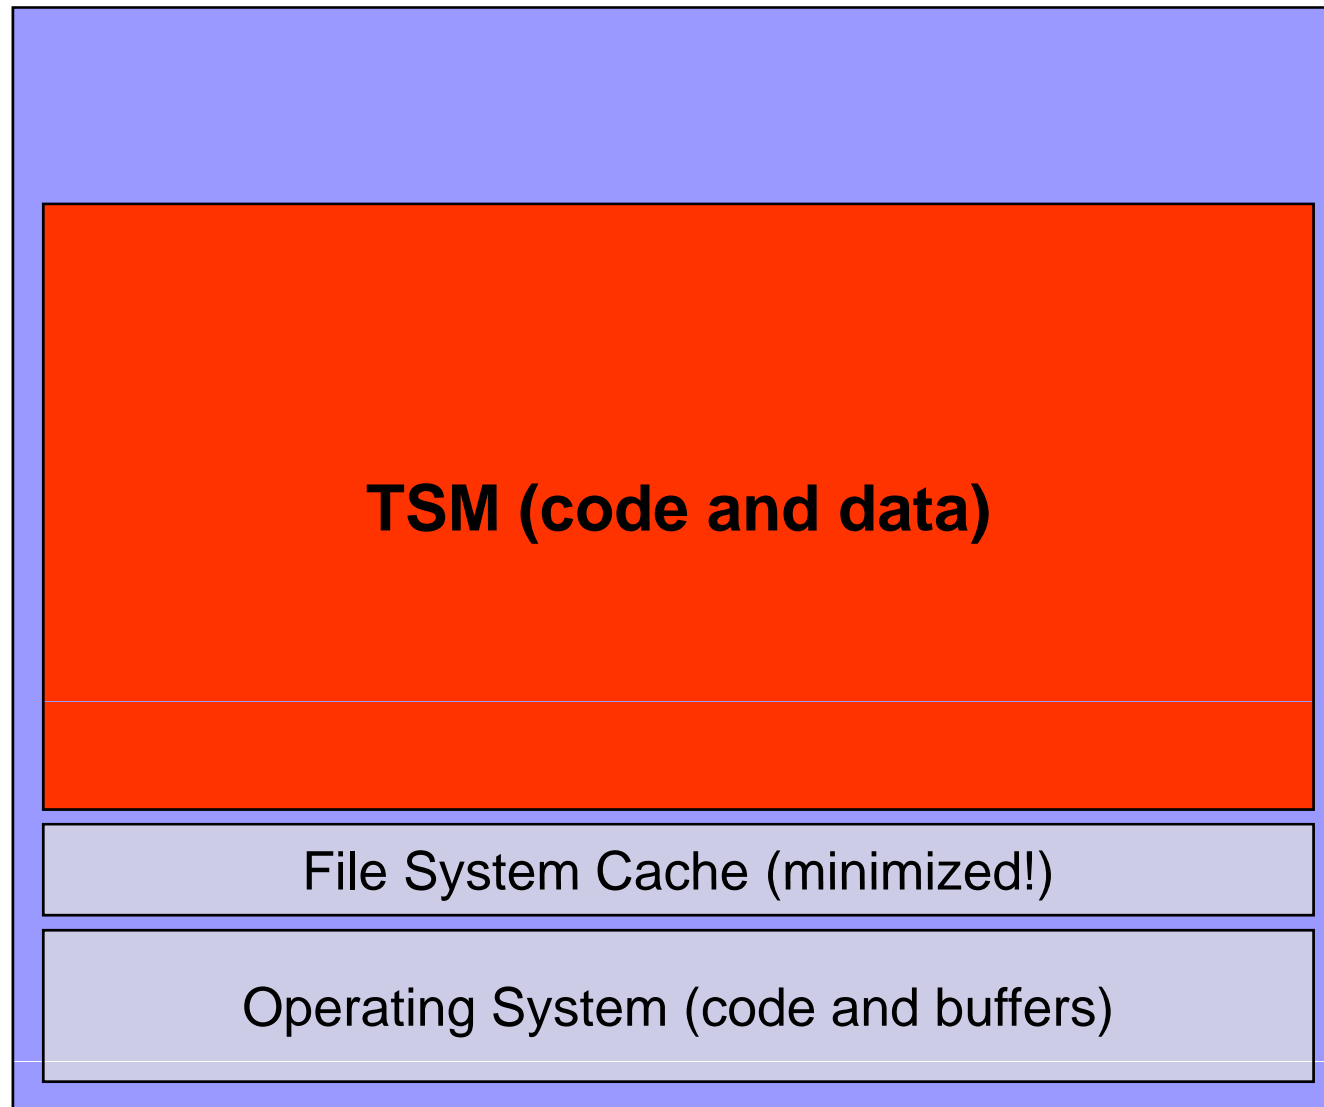
- Use all available memory
- Increase BUFPOOL as long as system does not start to page

- Simple to do
- May not fit your environment
- Only feasible for systems with small amount of RAM

- Use necessary memory
- Set BUFPOOLSIZE to $\frac{1}{2}$ of free memory
- Monitor „cache hit percentage“ with „query database format=detail“
- Increase memory if value < 98%

- Best way to do it but time-consuming

Memory layout



Memory analysis




- Use OS specific tools
- Use OS specific knowledge
- Pitfalls during analysis
 - Shared elements
 - Device buffers
 - File cache

- Samples taken on real customer TSM servers
- Samples taken on AIX with „nmon“
- Charts prepared with nmon analyzer
- Long term reporting on daily or weekly basis
- Periodic checks recommended
- Your mileage may vary

Increase BUFPOOL




Memory Use tsm1par1p2 08.07.2007



No paging




Paging tsm1par1p2 (pgspace) 08.07.2007



File IO




Paging tsm1par1p2 (filesystem) 08.07.2007



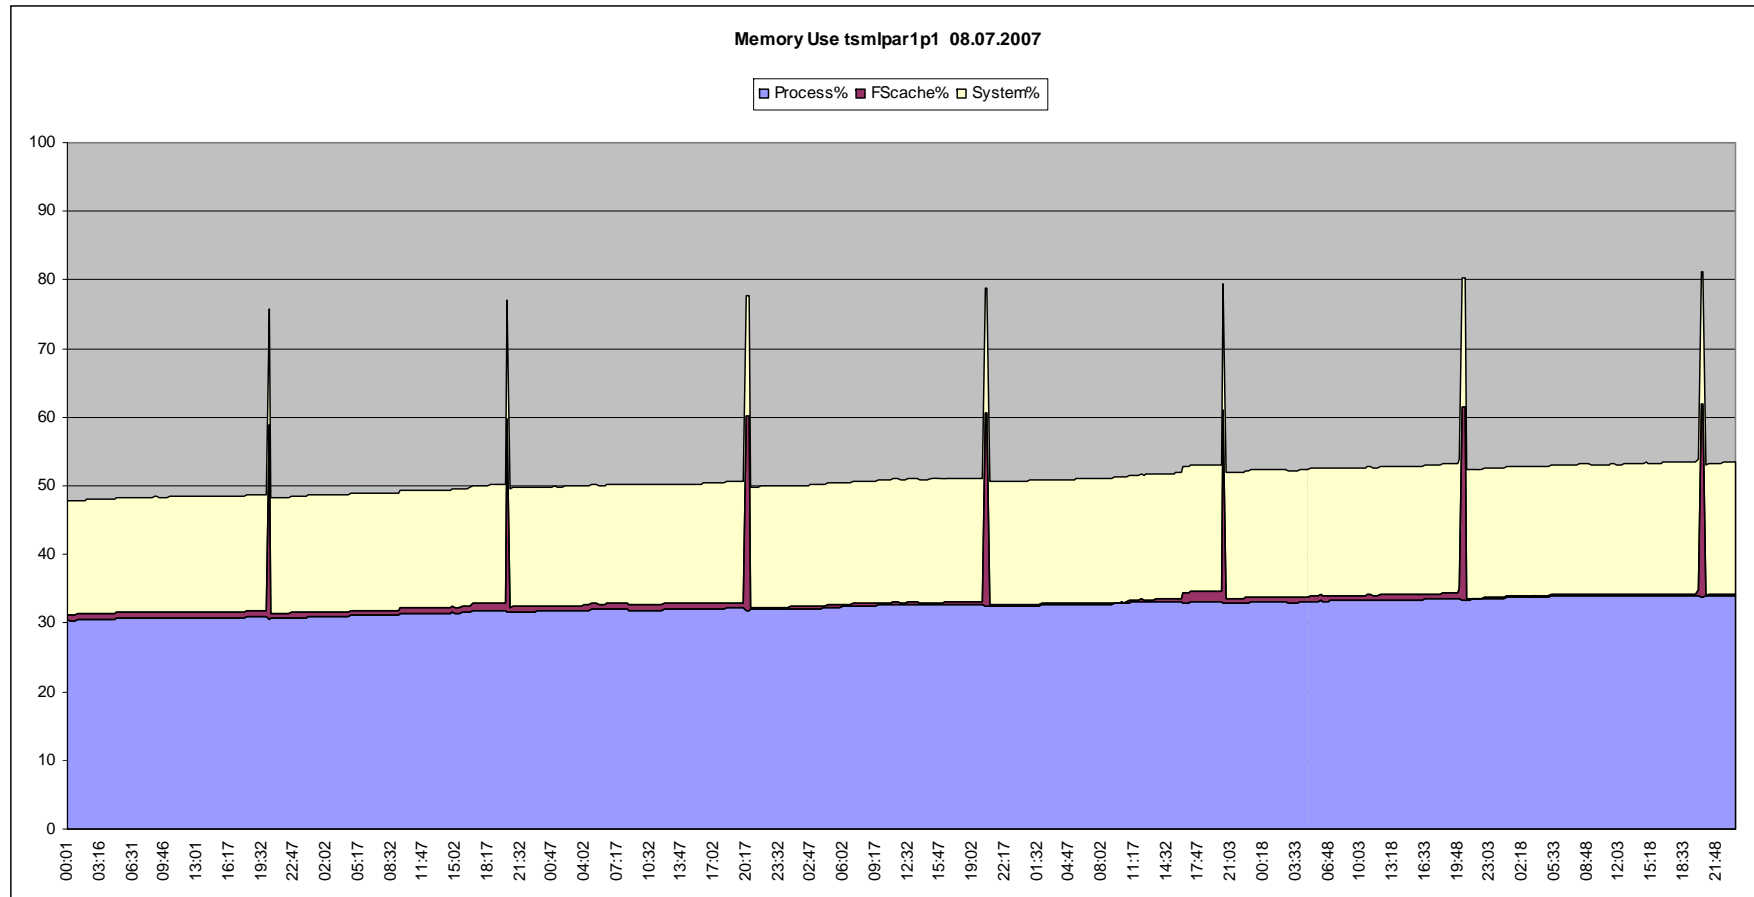
Process memory




Memory by command (MBytes) nsm4fix 16.07.2007



Daily incremental



Optimized memory



- What we know by now:
 - What is memory used for?
 - Does size really matter?
 - How to tune memory inside TSM
 - How to tune memory on the OS
 - How to get more out of less



Please contact



- address:
eXstor GmbH
Hans-Bredow-Strasse 40
D-65189 Wiesbaden
Germany
- E-Mail: info@exstor.de
- Web: www.exstor.de
- Phone: +49 611 724286-0