

Experiences with Oracle SPARC S7-2 Server

Marcel Hofstetter

hofstetter@jomasoft.ch

<https://jomasoftmarcel.blogspot.ch>

CEO / Enterprise Consultant
JomaSoft GmbH



Oracle ACE „Solaris“

Agenda

- About JomaSoft
- SPARC S7 CPU & Silicon Secured Memory & DAX
- SPARC S7-2 Server
- SPARC S7-2 Server @ JomaSoft
- SPARC S7 Project

About JomaSoft

- Engineering company founded in July 2000
- specialized in **Solaris** and software development, operations and consulting
- Product **VDCF** (Virtual Datacenter Cloud Framework)
Installation, Management, Operations, Monitoring, Security and DR for Solaris 10/11,
Virtualize using LDoms and Solaris Zones
- VDCF is used in production since 2006



Specialized
Oracle Solaris 11



Specialized
SPARC T-Series Servers



Bundesministerium
für Verkehr,
Innovation und Technologie



Marcel Hofstetter

Working in IT since 25+ years
Solaris since 20 years
CEO at JomaSoft GmbH since 18 years

International Speaker:
Oracle OpenWorld, DOAG, UKOUG, SOUG, AOUG



 **Oracle ACE „Solaris“**

SOUG (Swiss Oracle User Group) – Speaker of the Year 2016


Hobbies: Family, Travel, Wine & Dine, Movies

 <https://www.linkedin.com/in/marcelhofstetter>

 https://twitter.com/marcel_jomasoft

 <https://jomasoftmarcel.blogspot.ch>

Oracle SPARC CPUs compared



	SPARC M8 (2017)	SPARC S7 (2016)	SPARC M7 (2015)	SPARC T5 (2013)
Processor Cores	32 (5th Gen)	8 (4th Gen)	32 (4th Gen)	16 (3rd Gen)
Cache per Core	2 MB	2 MB	2 MB	0.5 MB
Memory Bandwidth per Core	5.6 GB/sec	6.0 GB/sec	5.3 GB/sec	5.0 GB/sec
Memory Access	127ns	97ns	131ns	163ns
I/O Bandwidth	145 GB/sec	32 GB/sec	145 GB/sec	32 GB/sec
CPU Frequency	5.0 GHz	4.27 GHz	4.13 GHz	3.6 GHz


Oracle SPARC M7/8 & S7 CPU



Security in Silicon:
Silicon Secured Memory
Cryptography Acceleration

SQL in Silicon:
Database In Memory Accelerator Engines

Capacity in Silicon:
Decompression Engines



SPARC – Silicon Secured Memory

- Integrated in SPARC CPU M7/M8 and S7
- You detect and prevent
 - Memory Reference Errors
 - Buffer Overruns
 - Memory Usage after free
- Alternatives in Software are expensive and 30x – 70x slower
- Oracle Developer Studio Compiler includes Support for Discovery at Development
- Demo Video about OpenSSL Heartbleed

https://swisdev.oracle.com/_files/ADI-Demo.html

SPARC – Silicon Secured Memory

```
void main(int argc, char *argv[])
{
    char *buffer = malloc( sizeof(char) * 10);
    strcpy(buffer, "Test-Text");
    for (int i = 0; i < 20; ++i)
        printf( "%c ", buffer[i] );
    printf("|\\n");
    free(buffer);
}
```

```
/opt/solarisstudio12.4/bin/cc -m64 -g -o buffer_overrun buffer_overrun.c
```

T	E	S	T	-	T	E	X	T			?		P	W	D				
---	---	---	---	---	---	---	---	---	--	--	---	--	---	---	---	--	--	--	--

```
-bash-4.4$ ./buffer_overrun
```

```
T e s t - T e x t |
```


SPARC – Silicon Secured Memory

With SSM (ADI) activated the Program is stopped and can't access foreign Memory.

```
-bash-4.4$ LD_PRELOAD_64=/lib/64/libadimalloc.so.1 ./buffer_overrun  
Segmentation Fault (core dumped)
```

```
-bash-4.4$ echo ::status | mdb core  
debugging core file of buffer_overrun (64-bit) from g0072  
file: /export/home/marcel/buffer_overrun  
initial argv: ./buffer_overrun  
threading model: native threads  
status: process terminated by SIGSEGV (Segmentation Fault), pc=100000bb0  
, ADI version d mismatch for VA ffffffff7e93ffc0
```

SPARC – Silicon Secured Memory

Detailed Results when using the Compiler Libraries

```
LD_PRELOAD_64=/opt/developerstudio12.5/lib/compiler/sparcv9/libdiscoverADI.so ./
```

```
buffer_overrun
```

```
T e s t - T e x t
```

1. ABR: reading memory beyond array bounds at address 0x2ffffff7cc7e040

```
main() + 0x60 (line ~12) in "buffer_overrun.c"
9: strcpy(buffer, "Test-Text");
10:
11: for (int i = 0; i < 20; ++i)
12: printf("%c ", buffer[i]);
13: printf("\n");
14:
15: free(buffer);

was allocated at (0x2ffffff7cc7e000, 64 bytes):

main() + 0x10 (line ~7) in "buffer_overrun.c"
4:
5: void main(int argc, char *argv[])
6: {
7: char *buffer = malloc( sizeof(char) * 10);
8:
9: strcpy(buffer, "Test-Text");
10:
```

INMEMORY / DAX


Test Setup using SLOB

```
SQL> show parameter inmemory_size
```

NAME	TYPE	VALUE
inmemory_size	big integer	1G

```
SQL> ALTER TABLE USER1.CF1 INMEMORY;
Table altered.
```

```
SQL> select count(*) from USER1.CF1;
COUNT(*)
-----
10000
```



INMEMORY / DAX

Result / 8 Reader / 1 x SPARC S7-core

awr_0w_8r.20181107_165153.txt

DB Name	DB Id	Unique Name	DB Role	Edition	Release	RAC	CDB
SLOB	3718155087	SLOB	PRIMARY	EE	18.0.0.0.0	NO	NO

Host Name	Platform	CPUs	Cores	Sockets	Memory (GB)
v0133	Solaris[tm] OE (64-bit)	8	1	1	16.00

	Snap Id	Snap Time	Sessions	Curs/Sess
Begin Snap:	105	07-Nov-18 16:46:32	44	1.3
End Snap:	106	07-Nov-18 16:51:51	44	1.3
Elapsed:		5.31 (mins)		
DB Time:		42.12 (mins)		

Load Profile	Per Second	Per Transaction	Per Exec	Per Call
DB Time (s):	7.9	126.4	0.00	8.15
DB CPU(s):	7.9	125.4	0.00	8.09
Background CPU(s):	0.0	0.5	0.00	0.00
Redo size (bytes):	8,690.5	138,454.2		
Logical read (blocks):	125,562,125.4	2,000,411,835.5		
Block changes:	42.8	681.1		
Physical read (blocks):	0.6	10.1		
Physical write (blocks):	3.2	51.2		
Read IO requests:	0.3	5.2		
Write IO requests:	1.4	22.8		
Read IO (MB):	0.0	0.1		
Write IO (MB):	0.0	0.4		
IM scan rows:	125,536,275.3	2,000,000,000.0		
Session Logical Read IM:	125,536,275.3	2,000,000,000.0		
User calls:	1.0	15.5		

-bash-4.4\$ grep offload awr_0w_8r.20181107_165153.txt

Statistic	Total	per Second	per Trans
IM simd compare HW offload calls	4,000,000	12,553.6	200,000.0
IM simd decode unpack HW offload	4,000,000	12,553.6	200,000.0

Result without INMEMORY / DAX

Result / 8 Reader / 1 x SPARC S7-core

awr_0w_8r.20181119_153421.txt

DB Name	DB Id	Unique Name	DB Role	Edition	Release	RAC	CDB
SLOB	3718155087	SLOB	PRIMARY	EE	18.0.0.0.0	NO	NO

Host Name	Platform	CPUs	Cores	Sockets	Memory(GB)
v0133	Solaris[tm] OE (64-bit)	8	1	1	16.00

	Snap Id	Snap Time	Sessions	Curs/Sess
Begin Snap:	277	19-Nov-18 15:21:56	49	.9
End Snap:	278	19-Nov-18 15:34:19	49	1.0
Elapsed:		12.38 (mins)		
DB Time:		98.54 (mins)		

Load Profile	Per Second	Per Transaction	Per Exec	Per Call
DB Time(s):	8.0	197.1	0.00	18.95
DB CPU(s):	7.9	195.3	0.00	18.78
Background CPU(s):	0.0	0.6	0.00	0.00
Redo size (bytes):	4,907.6	121,553.6		
Logical read (blocks):	1,386,001.4	34,328,851.4		
Block changes:	22.2	548.7		
Physical read (blocks):	27.3	675.		
Physical write (blocks):	1.2	28.9		
Read IO requests:	26.7	662.2		
Write IO requests:	0.6	13.6		

1.3 Mio read anstatt 125 Mio read

INMEMORY / DAX

The 1 core LDOM uses all 4 DAX Units of the SPARC S7 Socket.

Host Name	Platform	CPUs	Cores	Sockets	Memory (GB)
v0133	Solaris[tm] OE (64-bit)	8	1	1	16.00

```
-bash-4.4$ daxstat 10
```


DAX	commands	fallbacks	input	output	%busy
4	63809	0	106.1M	5.4M	0
5	63810	0	106.1M	5.5M	0
6	63810	0	106.1M	5.4M	0
7	63803	0	106.1M	5.5M	0

INMEMORY / DAX



8 core LDOM - compressed and uncompressed compared

compressed (low) 130 MB

uncompressed 560 MB



Oracle SPARC S7 Server



	S7-2	S7-2L
CPU	1-2	2
Cores	8/16	16
Threads	64/128	128
Max Memory	1 TB	1 TB
Size	1 RU	2 RU
PCI Slots	3 x PCIe 3.0	6 x PCIe 3.0


Oracle SPARC S7-2 Server



	S7-2 small	S7-2 medium
CPU	2	2
Cores	16	16
Disk	2 x 600 GB	2 x 600 GB
Memory	128 GB	512 GB
Price	ca. 11K £	ca. 19K £

SPARC – Single Thread Performance

Single Thread Performance Compare (execution time in seconds)



Benchmark done by Oracle

SPECjEnterprise2010 (Java App & DB Server)

SPARC S7-2 / 512 GB

16 cores 14,400 EjOPS 900 OPS per Core

Oracle X6-2 / 256 GB

44 cores 27'800 EjOPS 631 OPS per Core

SPARC S7 43% better result at core to core compare

Details:

<https://blogs.oracle.com/bestperf/specjenterprise2010:-sparc-s7-2-secure-and-unsecure-results>

<https://www.spec.org/jEnterprise2010/results/jEnterprise2010.html>

Benchmark done by Oracle

SPECjEnterprise2010 (Java App & DB Server)

SPARC S7-2 / 512 GB

16 cores

19K £

5100 JOPS

900 OPS per Core

Oracle X6-2 / 256 GB

44 cores

14K £

5100 JOPS

631 OPS per Core

SPARC S7 43% better result at core to core compare

Details:

<https://blogs.oracle.com/bestper>

<https://www.spec.org/jEnterprise2010/resu>

Think about App & DB Licenses
Less core, less cost

JomaSoft Development T4 → S7

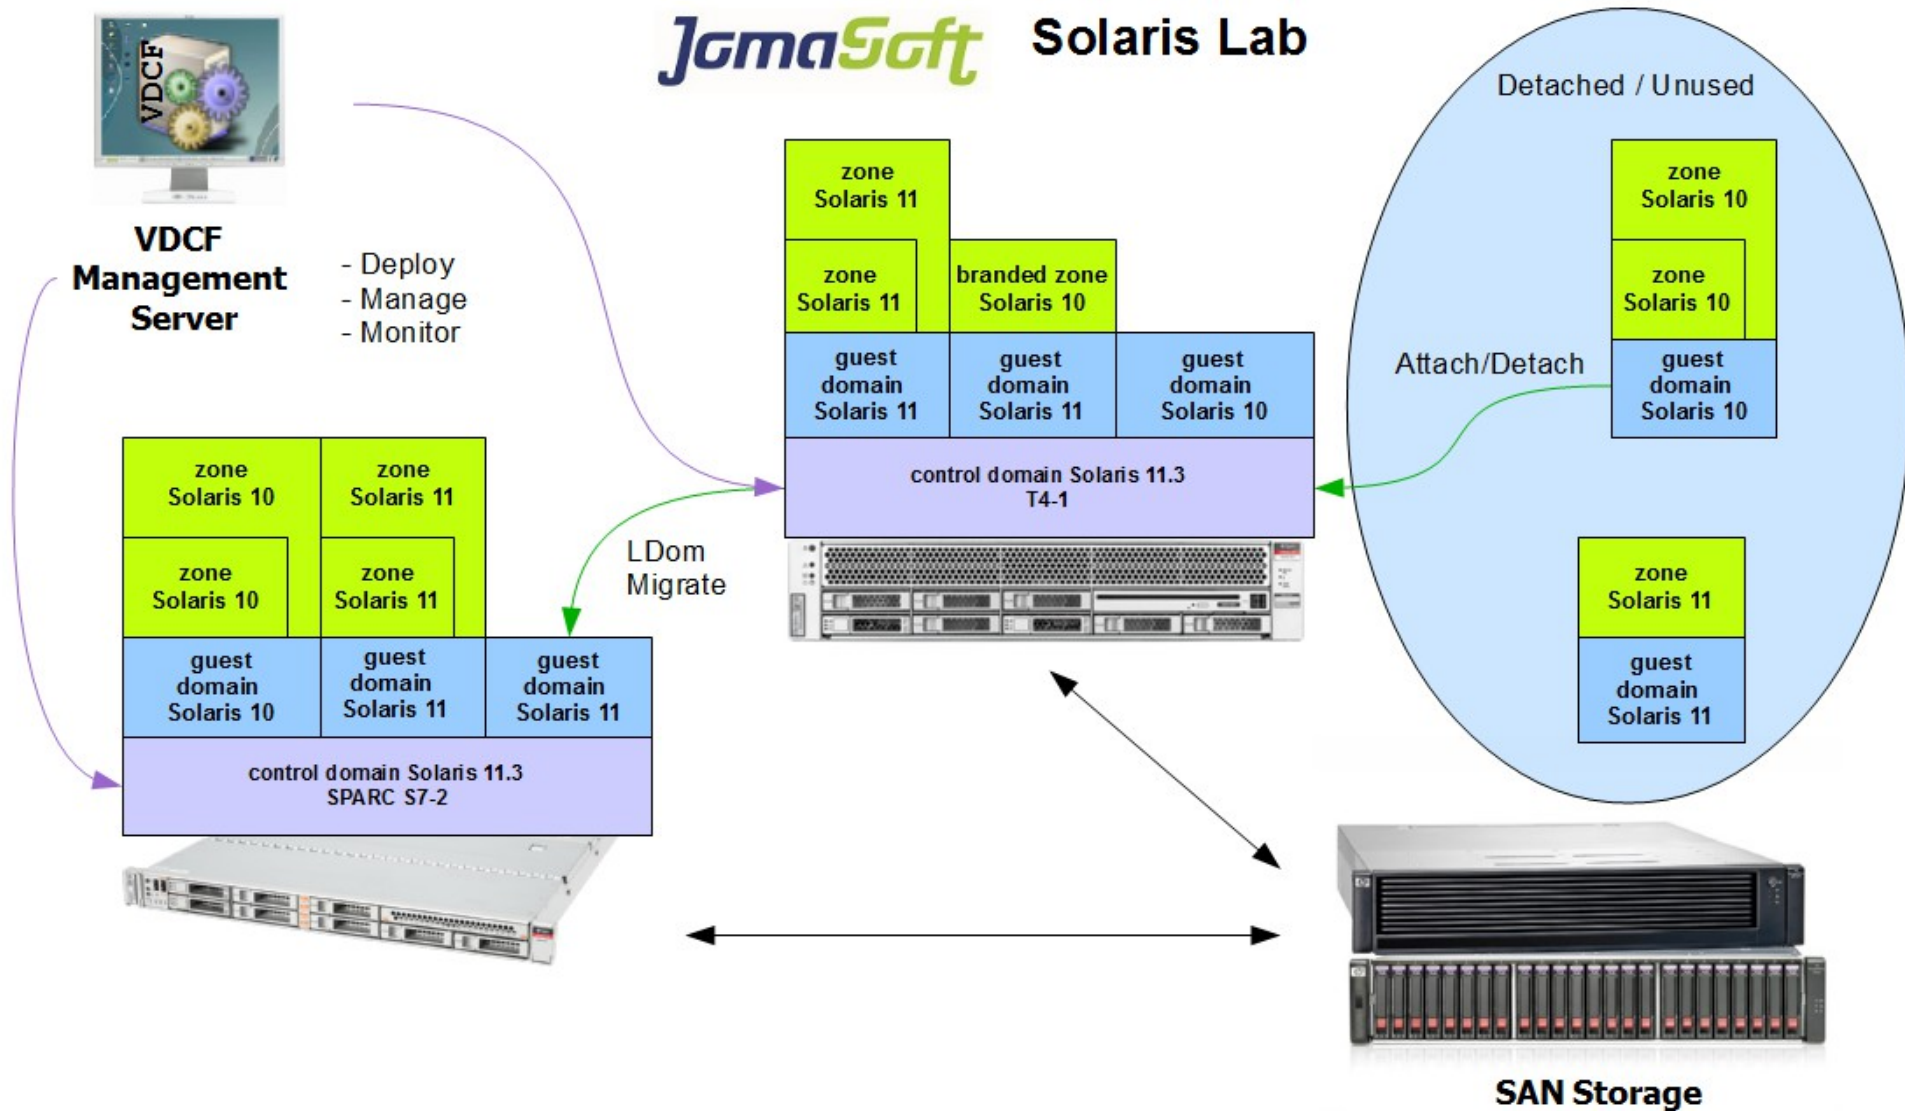
7 years old, works like a charm



	T4-1	S7-2
CPU	1	2
Cores	8	16
Freq	2.85 GHz	4.27 GHz
Memory	32 GB	128 GB

- VDCF Application performance: 2x
(Ldom with 1 Core used)

JomaSoft Development & Test



JomaSoft Development & Test

[Home](#) -
 [CDoms](#)
[Nodes](#)
[vServers](#)
[Compliance Reports](#) -
 [logout](#)

Guest Domain list

Show entries

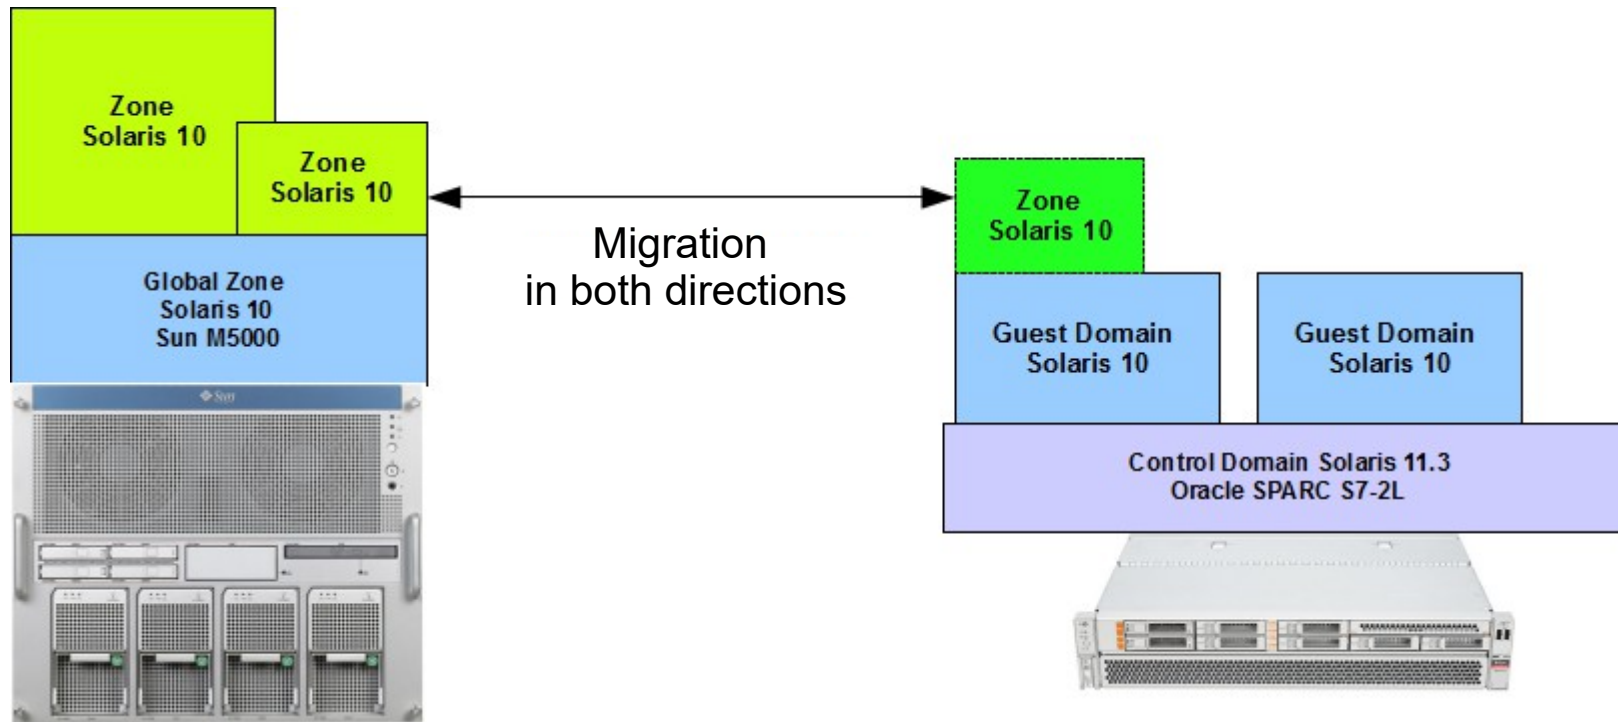
Search:

Name	cState	rState	CDom	cPool	Model	Build	CPU Cores	CPU Max Cores	vCPUs	Memory (GB)	vServers	Comment	Compliance Score
g0040	ACTIVE	ACTIVE (RUNNING)	s0003	ist-stand	ORCL,SPARC-S7-2	5.10svz_u11_req	0	0	4	4.0	1	s10 ipmp	
g0044	ACTIVE	ACTIVE (RUNNING)	s0003	sol11	ORCL,SPARC-S7-2	s11u3-sru27	1	0	8	16.0	0	Marcel DB Perf Testing SLOB	
g0049	ACTIVE	ACTIVE (RUNNING)	s0003	sol11	ORCL,SPARC-S7-2	s11u4-sru1	0	0	4	4.0	1	Marcel S11.4	
g0056	DETACHED	-	(s0003)	sol11	ORCL,SPARC-S7-2	s11u3-sru14-uar	0	0	2	4.0	3	MECH Puppet (Agent)	90.2
g0059	ACTIVE	ACTIVE (RUNNING)	s0003	sol11	ORCL,SPARC-S7-2	s11u4-sru2	1	0	8	16.0	1	SLOB mit 18c	
g0062	ACTIVE	ACTIVE (RUNNING)	s0003	ist-stand	ORCL,SPARC-S7-2	s11u3-sru35	0	0	8	8.0	3	ZFS cloning / Shared DS	94.8
g0081	ACTIVE	ACTIVE (RUNNING)	s0003	sol11	ORCL,SPARC-S7-2	s11u3-sru36	1	0	8	4.0	1	LDom Autotest	
g0086	DETACHED	-	(s0003)	sol11	ORCL,SPARC-S7-2	s11u4-sru2	1	0	8	4.0	0	Solaris 11 - Desktop	100.0
g0092	ACTIVE	ACTIVE (RUNNING)	s0003	sol11	ORCL,SPARC-S7-2	s11u3-sru35	0	0	1	3.0	3	MECH Upgrade	

Showing 1 to 9 of 9 entries

[Previous](#)
[1](#)
[Next](#)

Hardware Life Cycle Project



Perfect starting point, because Customer already used Zones

Done : - Setup of 2 x S7-2L
- Zones Migration from 4 x M5000 Domains into 4 new LDoms

Duration: 6 days

Experiences with Oracle SPARC S7-2 Server

Questions?

Marcel Hofstetter

hofstetter@jomasoft.ch

<https://jomasoftmarcel.blogspot.ch>

CEO / Enterprise Consultant
JomaSoft GmbH



Oracle ACE „Solaris“

 <https://www.linkedin.com/in/marcelhofstetter>

 https://twitter.com/marcel_jomasoft

 <https://jomasoftmarcel.blogspot.ch>

My second Session at #UKOUG_TECH18

What is Delivered with Solaris 11.4?

Wednesday, 5.12. 11:40