

The Power-per-Core Dilemma

Dierk Lenz

DOAG Konferenz

15. November 2016



- Est. 1996
- Located in Burscheid (near Leverkusen) in Germany
- Consulting, Courses, Operations for Oracle Databases
- High Availability, Tuning, Migrations and Troubleshooting
Monitoring
- Booth at Level 2 (236)
 - Products: Monitoring Module, Taskzone

The Dilemma

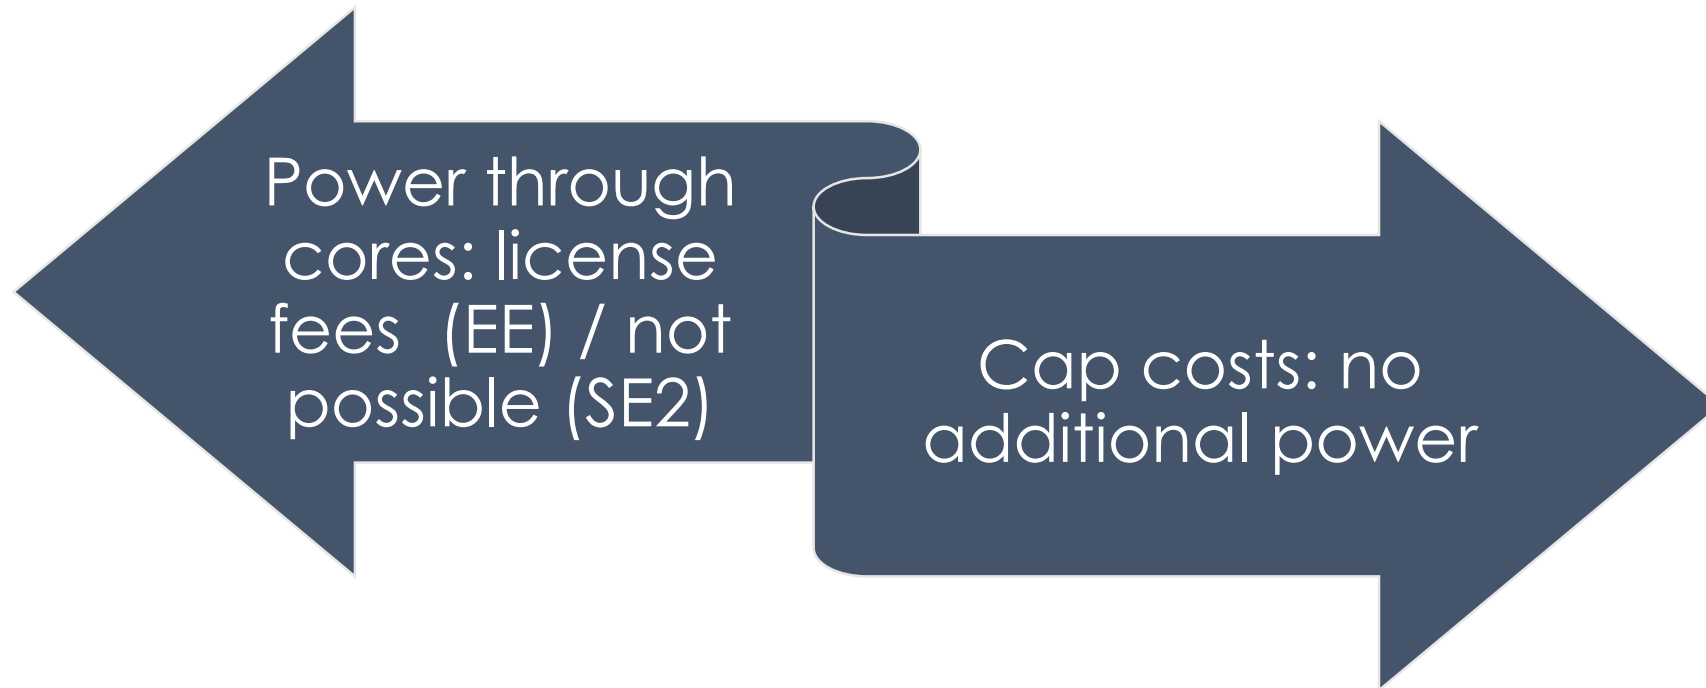
Dilemma

A situation, in which one has to choose between two bad possibilities.

(Source: Google, Translation: Me)

Prerequisites

- More power for modern computer systems through more cores
- Cores do not cost much, but:
 - License fees for Oracle Database Enterprise Edition (EE) linear dependent of the number of cores
 - Oracle Database Standard Edition Two (SE2) limits the number of used cores



Oracle Database Enterprise Edition

- „Every core costs“
 - Calculation of processors for licensing
 $\#cores * \langle \text{core factor} \rangle$
i.e. for Intel/AMD processors:
 $18 \text{ cores} * 0,5 = \mathbf{9 \text{ processors}}$
- Licensing with Named User Plus defines a minimum of 25 NUP per processor

Oracle Database Standard Edition 2

- „Old“ SE/SE1 licenses no longer supported in some months*
- Usage of SE2 limited to systems with a maximum of 2 sockets
- Only 16 cpu threads used at the same time per instance

*Fees for Extended Support waived until May 2017

Virtualization

Virtualization

- Used for numerous systems
- Oracle licensing distinguishes between soft and hard partitioning
 - Most Virtualization solutions (e.g. VMWare): soft partitioning
 - Exceptions Solaris, AIX, HP-UX and Oracle VM (configured with processor pinning)
 - Soft partitioning not valid for limitation of cores to license
- Complete servers to license in most cases!

A Risk: Live Migration

- Live Migration (or similar features) moves VMs online between servers
- New VMWare versions Live Migration possible without limits
- Licensing of full environment
 - In case some servers are not SE2-capable (more than two CPUs): EE licenses for everything!



Recommendation: Avoid Virtualization

- Risks of virtualization
 - High risk for additional license fees
 - High complexity
 - Performance drop

Processors



Recommendation: few strong cores instead of lots of weak cores !

- EE licensing not dependent on strength of cores
- 16 CPU thread limit:

Be as efficient as possible when calculating DB tasks

- Number of cores is not the complete story!
- Take processor types and core factors into account!
(in most cases:
 - Intel/AMD ~ 0,5
 - PowerPC ~ 1
 - SPARC ~ 0,5)

Possibilities for Growth?

- Limiting the hardware „in use“ is often not reflected by the number of licenses needed
- Exceptions for hard partitioning
- Big advantage when using Oracle Engineered Systems
 - *Capacity On Demand* (available on most systems):
Just use the licensed number of cores!
(without the need for Oracle VM)
 - *Trusted Partitions* (Oracle Private Cloud Appliance):
Use the licensed number of cores under Oracle VM without processor pinning

Storage

Datenbases and SAN/NAS

- Database performance extremely dependent on IO performance (throughput and latency!)
- Commonly used systems may cause problems
- No problems when using local/private devices!

Flash and SSD

- Current situation:
 - Better quality and performance
 - Lower prices
- Getting better and better for databases!
- Caution: limited lifetime!

Recommendations

- Use local storage (not shared)
- Use Flash/SSD

Oracle Engineered Systems

E.G. Oracle Database Appliance X6-2

- New series of ODAs
- Fast installation and configuration
- Patching of full software stack
- CPUs with 10 cores
- Flash/SSD

Komponente	ODA X6-2S	ODA X6-2M	ODA X6-2L	ODA X5-2	ODA X6-2HA
Knoten / Server	1 x	1 x	1 x	2 x	2 x
CPU	1 Socket x 10 Kerne	2 Socket x 10 Kerne	2 Socket x 10 Kerne	2 Socket x 18 Kerne pro Server	2 Socket x 10 Kerne pro Server
Memory	128 GB (384 GB max.)	256 GB (768 GB max.)	256 GB (768 GB max.)	256 GB (512 GB max.) pro Server	256 GB (768 GB max.) pro Server
Storage (Daten)	NVMe 6,4 TB raw: 3,2 TB netto 12,8 TB raw: 6,4 TB / 4,2 TB netto	NVMe 6,4 TB raw: 3,2 TB netto 12,8 TB raw: 6,4 TB / 4,2 TB netto	NVMe 19,2 TB raw: 9,6 TB / 6,4 TB netto 28,2 TB raw: 14,4 TB / 9,6 TB netto	SAS 128 TB raw (256 TB max.) 64 TB netto gespiegelt 42,7 TB dreifach gespiegelt + SSD 4 x 400 GB + SSD 4 x 200 GB	SSD (Storage Shelf) 12 TB raw: 6 TB / 4 TB netto 24 TB raw: 12 TB / 8 TB netto 48 TB raw: 24 TB / 16 TB netto + extra SSD's für Redo Log
Boot Disks	SSD 480 GB gespiegelt	SSD 480 GB gespiegelt	SSD 480 GB gespiegelt	SAS 600 GB gespiegelt pro Server	SSD 480 GB gespiegelt pro Server
Netzwerk	2 x 10 Gb 2 x FC	4 x 10 Gb 2 x FC	4 x 10 Gb 2 x FC	4 x 10 Gb Infiniband Interconnect, FC optional	4 x 10 Gb Infiniband Interconnect, FC optional
Capacity on	Ja – 2, 4, 6, 8, 10 Kerne	Ja – 2, 4, 6, 8, 10, 12, 14, 16, 18, 20 Kerne	Ja – 2, 4, 6, 8, 10, 12, 14, 16, 18, 20 Kerne	Ja – 2, 4, 6, 8, bis 36 Kerne pro Server	Ja – 2, 4, 6, 8, bis 20 Kerne pro Server
Enterprise Edition	Ja	Ja	Ja	Ja	Ja
Standard Edition	Ja	Ja	Ja	Nein	Nein
Virtualisierung (OracleVM)	Nein	Nein	Nein	Ja	Ja
RAC	Nein	Nein	Nein	Ja	Ja

High Availability

Standby Databases

- Standby databases more important for big databases
- Especially when using Flash/SSD
- In case of crash: failover to standby instead of restoring a backup
- For **Enterprise Edition:**

Data Guard!

- Proven for **Standard Edition:**

Dbvisit Standby

- New version 8 with a bunch of new features

Monitoring



```
top - 14:56:08 up 4 days, 3:38, 3 users, load average: 1.40, 1.34, 1.54
Tasks: 310 total, 2 running, 308 sleeping, 0 stopped, 0 zombie
Cpu(s): 25.3%us, 0.8%sy, 0.0%ni, 73.6%id, 0.2%wa, 0.0%hi, 0.0%si, 0.0%st
Mem: 32880824k total, 32714984k used, 165840k free, 10612k buffers
Swap: 4321276k total, 386772k used, 3934504k free, 13426176k cached
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
27729	oracle	20	0	15.9g	1.4g	29m	R	97.8	4.4	32:23.82	ora_j000- [REDACTED]
1925	root	20	0	53612	1832	1528	S	1.0	0.0	4:31.33	vmtoolsd
8373	oracle	-2	0	14.5g	16m	15m	S	1.0	0.1	74:25.48	ora_vktn- [REDACTED]
17	root	20	0	0	0	0	S	0.3	0.0	0:06.10	ksoftirqd/3
4570	[REDACTED]	20	0	15160	1372	924	S	0.3	0.0	1:33.12	top
16289	root	20	0	4132	1072	524	S	0.3	0.0	0:22.77	nailsd
16794	nails	20	0	243m	13m	2112	S	0.3	0.0	1:23.00	scanner
1	root	20	0	19232	1044	836	S	0.0	0.0	0:10.09	init
2	root	20	0	0	0	0	S	0.0	0.0	0:00.47	kthreadd
3	root	RT	0	0	0	0	S	0.0	0.0	0:06.92	migration/0
4	root	20	0	0	0	0	S	0.0	0.0	0:08.89	ksoftirqd/0
5	root	RT	0	0	0	0	S	0.0	0.0	0:00.00	stopper/0
6	root	RT	0	0	0	0	S	0.0	0.0	0:00.55	watchdog/0
7	root	RT	0	0	0	0	S	0.0	0.0	0:12.45	migration/1
8	root	RT	0	0	0	0	S	0.0	0.0	0:00.00	stopper/1
9	root	20	0	0	0	0	S	0.0	0.0	0:07.54	ksoftirqd/1
10	root	RT	0	0	0	0	S	0.0	0.0	0:00.82	watchdog/1

Monitoring

- Keep an eye on important values
- Emails for warnings and critical situations
- Fast reaction
- Historical data
- Performance and problem analysis
- HL Monitoring Module!

Applications

Needed: Efficient Applications!

- Avoid the Framework worst case scenario!
- *Designed for Performance!*
(not: first thoughts on performance when it's getting slow...)
- Keep an eye on well known facts for efficient database programming
(not: what are the new performance parameters for version X?)


Current Situation

- When do we know of performance problems?
 - As soon as the users complain!
 - When the system administrator detects higher system load!
 - Logging: All or Nothing
 - ➔ difficult analysis

- ***Why not, when a SQL statement takes longer than 1 sec?***

Instrumenting the Application

- It's possible... with Apps!
- Identify problems and „report to server“
- Intelligent Applications!
- Applications with performance intelligence!



More Presentations

Zeit	Ort	Vortrag
16.11. 11:00	St. Petersburg	The Battle: Linux vs. Windows Johannes Ahrends/Dierk Lenz Moderator: Martin Klier
16.11. 13:00	Shanghai	Zünde den Turbo-Boost – LOB-Migration beschleunigt Susanne Jahr
16.11. 15:00	Konferenzraum EG	Nachhaltiges Monitoring von Oracle-Datenbanken Rastislav Ciganek, Sascha Westermann
17.11. 09:00	Prag	Alles ist eine Aufgabe – mit Struktur erfolgreich sein Peter Bekiesch



Johannes Ahrendt, Dierk Lenz, Patrick Schwank, Günter Unbrusch

**Oracle 11g
Release 2
für den DBA**

Produktive Umgebungen effizient
konfigurieren, optimieren und verwalten

 ADDISON-WESLEY





Thank you for
your attention!